

Non-commutative Edmonds' problem and matrix semi-invariants

Gábor Ivanyos *

Youming Qiao †

K. V. Subrahmanyam ‡

June 20, 2016

Abstract

In 1967, J. Edmonds introduced the problem of computing the rank over the rational function field of an $n \times n$ matrix T with integral homogeneous linear polynomials. In this paper, we consider the *non-commutative version of Edmonds' problem*: compute the rank of T over the free skew field. This problem has been proposed, sometimes in disguise, from several different perspectives, in the study of e.g. the free skew field itself (Cohn 1973), matrix spaces of low rank (Fortin-Reutenauer, 2004), Edmonds' original problem (Gurvits, 2004), and more recently, non-commutative arithmetic circuits with divisions (Hrubeš and Wigderson, 2014).

It is known that this problem relates to the following invariant ring, which we call the *\mathbb{F} -algebra of matrix semi-invariants*, denoted as $R(n, m)$. For a field \mathbb{F} , it is the ring of invariant polynomials for the action of $\mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ on tuples of matrices – $(A, C) \in \mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ sends $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ to $(AB_1C^T, \dots, AB_mC^T)$. Then those T with non-commutative rank $< n$ correspond to those points in the nullcone of $R(n, m)$. In particular, if the nullcone of $R(n, m)$ is defined by elements of degree $\leq \sigma$, then there follows a $\mathrm{poly}(n, \sigma)$ -time randomized algorithm to decide whether the non-commutative rank of T is full. To our knowledge, previously the best bound for σ was $O(n^2 \cdot 4^{n^2})$ over algebraically closed fields of characteristic 0 (Derksen, 2001).

We now state the main contributions of this paper:

- We observe that by using an algorithm of Gurvits, and assuming the above bound σ for $R(n, m)$ over \mathbb{Q} , deciding whether or not T has non-commutative rank $< n$ over \mathbb{Q} can be done *deterministically* in time polynomial in the input size and σ .
- When \mathbb{F} is large enough, we devise an algorithm for the non-commutative Edmonds problem in time polynomial in $(n+1)!$. Furthermore, due to the structure of this algorithm, we also have the following results.
 - If the commutative rank and the non-commutative rank of T differ by a constant, then there exists a randomized efficient algorithm to compute the non-commutative rank of T . This improves a result of Fortin and Reutenauer, who gave a randomized efficient algorithm to decide whether the commutative and non-commutative ranks are equal.
 - We show that $\sigma \leq (n+1)!$. This not only improves the bound obtained from Derksen's work over algebraically closed field of characteristic 0 but, more importantly, also provides for the first time an explicit bound on σ for matrix semi-invariants over fields of positive characteristics. Furthermore, this does not require \mathbb{F} to be algebraically closed.

2010 Mathematics Subject Classification: Primary 13A50, 68W30.

Keywords: Edmonds' problem, symbolic determinant identity test, semi-invariants of quivers, non-commutative rank

*Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary. E-mail: Gabor.Ivanyos@sztaki.mta.hu.

†Centre for Quantum Computation and Intelligent Systems, University of Technology Sydney, Australia. E-mail: jimmyqiao86@gmail.com.

‡Chennai Mathematical Institute, Chennai, India. E-mail: kv@cmi.ac.in.

1 Introduction

1.1 The non-commutative Edmonds problem

In 1967, J. Edmonds introduced the following problem [30]: let $X = \{x_1, \dots, x_m\}$ be a set of variables. Given an $n \times n$ matrix T whose entries are homogeneous linear polynomials from $\mathbb{Z}[X]$, determine the rank of T over the rational function field $\mathbb{Q}(X)$, denoted as $\text{rk}(T)$. The decision version of Edmonds' problem is to decide whether T is of full rank or not; this decision version is better known now as the symbolic determinant identity testing (SDIT) problem. It is natural to consider this problem over any field \mathbb{F} . If $|\mathbb{F}|$ is constant, this problem is NP-hard [9]. This is not the setting we are concerned with – we will always assume $|\mathbb{F}|$ to be at least $\Omega(n)$.

When $|\mathbb{F}| \geq 2n$, the Schwartz-Zippel lemma provides a randomized efficient algorithm. To devise a deterministic efficient algorithm has a long history, and is of fundamental importance in complexity theory. Originally, the main motivation was its applications to certain combinatorial problems, most notably the maximum matching problem on graphs, as exploited by Tutte [67], Edmonds [30], Lovász [52], among others.¹ Since 2003, a major incentive to study SDIT arises from its implications to circuit lower bounds, as shown in the wonderful work by Kabanets and Impagliazzo [48]. Improving the results in [48], Carmosino et al. showed that such an algorithm implies the existence of an explicit multilinear polynomial family such that its graph is computable in NE, but the polynomial family cannot be computed by polynomial-size arithmetic circuits [10].

In this paper, we study Edmonds' problem in the non-commutative setting. In other words, we view the entries of T as elements of $\mathbb{F}\langle X \rangle$, the algebra of non-commutative polynomials over \mathbb{F} . To state this, we need a non-commutative counterpart of the rational function field. Note that, due to non-commutativity, the best we can hope for is a skew field (a.k.a. a non-commutative field or a division ring). The *free skew field* is the non-commutative analogue of the rational function field. For T a matrix with homogeneous linear polynomials, Fortin and Reutenauer [33] defined the *non-commutative rank* of T , denoted as $\text{ncrk}(T)$, as its rank over the free skew field. By the *non-commutative Edmonds problem* we mean the problem of computing $\text{ncrk}(M)$, and by the *non-commutative full rank problem* (NCFullRank) we mean the problem of deciding whether $\text{ncrk}(M)$ is full or not.

Since we will not define the free skew field formally, we provide a definition of $\text{ncrk}(T)$ that is due to Cohn [16]: $\text{ncrk}(T)$ is the minimum $s \in \mathbb{Z}^+$ s.t. T can be written as PQ , where P and Q are matrices with homogeneous linear polynomials from $\mathbb{F}\langle X \rangle$, and of size $n \times s$ and $s \times n$, respectively. It may also be worthwhile to recall that, since every module over a skew field is free, the row (resp. column) rank of a matrix over a skew field can be defined as the rank of the module generated by the rows (resp. columns). Then just as the case of matrices over fields, it can be shown that row rank and column rank are equal, so either of them defines the rank of a matrix over a skew field.

The free skew field was first constructed by Amitsur [3], and alternative constructions were subsequently given by Bergman [5], Cohn [15], and Malcolmson [54]. We refer the reader to [43] by Hrubeš and Wigderson for a nice introduction to the free skew field from the perspective of algebraic computations. Cohn's books [15, 16] serve as a comprehensive introduction to this topic.

It will be clear soon that $\text{rk}(T) \leq \text{ncrk}(T)$, and Fortin and Reutenauer showed that $\text{ncrk}(T) \leq 2\text{rk}(T)$, and exhibited an example T for which $\text{ncrk}(T) = 3/2 \cdot \text{rk}(T)$ [33]. In [17], Cohn and Reutenauer presented an algorithm to decide whether $\text{ncrk}(T)$ is full or not², which puts this problem in PSPACE since it reduces to testing the solvability of a system of multivariate polynomial equations. Unlike its commutative counterpart, it is not even clear that the non-commutative Edmonds problem has a randomized efficient algorithm. In Section 1.3, we will discuss a natural randomized algorithm for NCFullRank, but its efficiency will depend on an invariant-theoretic quantity.

¹In these applications, T is usually of certain specific forms, for example, as a mixed matrix: each entry is a single variable or a field element, and each variable appears only once.

²As remarked in [33], this algorithm can be generalized to compute $\text{ncrk}(T)$.

1.2 Equivalent formulations of the non-commutative Edmonds problem

Like Edmonds' problem, its non-commutative counterpart also has a long history, though in the literature, it often stated in a very different way. In 1973, Cohn first studied this problem from the perspective of understanding the free skew field, and showed it to be decidable in [13, 14]. In 2003, Gurvits posed this problem in his remarkable work on Edmonds' problem [39]. Recently, Hrubeš and Wigderson arrived at this problem in their study of non-commutative arithmetic circuits with divisions [43]. Indeed, a very intriguing feature of the non-commutative Edmonds problem is the existence of several interesting equivalent formulations. Instead of relying on the free skew field, these formulations use either linear algebra, or concepts from invariant theory, or quantum information theory. They are scattered in the literature, so we collect them here, to illustrate the various facets of this problem, introduce some previous works, and motivate the study of the non-commutative Edmonds problem.

To state these formulations we need some notations. $M(n, \mathbb{F})$ denotes the linear space of $n \times n$ matrices over \mathbb{F} . A linear subspace of $M(n, \mathbb{F})$ is called a *matrix space*. Given T , a matrix of linear forms in variables $X = \{x_1, \dots, x_m\}$, write $T = x_1 B_1 + x_2 B_2 + \dots + x_m B_m$ where $B_i \in M(n, \mathbb{F})$. Let $\mathcal{B} := \langle B_1, \dots, B_m \rangle$, where $\langle \cdot \rangle$ denotes linear span. The rank of \mathcal{B} , denoted as $\text{rk}(\mathcal{B})$, is defined as $\max\{\text{rk}(B) \mid B \in \mathcal{B}\}$. We call \mathcal{B} *singular*, if $\text{rk}(\mathcal{B}) < n$. When $|\mathbb{F}| > n$, as we will assume throughout, $\text{rk}(T) = \text{rk}(\mathcal{B})$.³ We shall soon see that $\text{ncrk}(T)$ corresponds to some property of \mathcal{B} as well, so that we can translate the study of commutative and non-commutative ranks of T entirely to the study of \mathcal{B} .

Some of these formulations make sense only subject to certain conditions. In such cases we indicate the conditions needed before that formulation.

- Given $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, a subspace $U \leq \mathbb{F}^n$ is called a *c-shrunk subspace* of \mathcal{B} , for $c \in \mathbb{N}$, if there exists $W \leq \mathbb{F}^n$, such that $\dim(W) \leq \dim(U) - c$, and for every $B \in \mathcal{B}$, $B(U) \leq W$. U is called a shrunk subspace of \mathcal{B} , if it is a *c-shrunk subspace* for some $c \in \mathbb{Z}^+$.

Question: compute the maximum c such that there exists a *c-shrunk subspace*.

Remark: Cohn showed that the non-commutative rank is not full if and only if there is a shrunk subspace [16]. This was generalized by Fortin and Reutenauer [33, Theorem 1] who showed a precise relationship between non-commutative rank and the existence of *c-shrunk subspaces*. Their motivation to consider this problem was to connect matrices over linear forms on the one hand, and matrix spaces of low rank on the other. The latter topic was studied in e.g. [4, 31]. By [33], we can define the *non-commutative rank* of \mathcal{B} as

$$n - \max\{c \in \{0, 1, \dots, n\} \mid \exists \text{c-shrunk subspace of } \mathcal{B}\},$$

and it follows that $\text{ncrk}(\mathcal{B}) = \text{ncrk}(T)$. So we may (and do) identify T with \mathcal{B} in the following.

- (\mathbb{F} is large enough) Given $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, the d th tensor blow-up of \mathcal{B} , is $\mathcal{B}^{[d]} := M(d, \mathbb{F}) \otimes \mathcal{B} \leq M(dn, \mathbb{F})$. It is clear that $\text{rk}(\mathcal{B}^{[d]}) \geq d \cdot \text{rk}(\mathcal{B})$. We shall prove that when \mathbb{F} is large enough, then d always divides $\text{rk}(\mathcal{B}^{[d]})$. Furthermore, when $d > n$, then $\text{rk}(\mathcal{B}^{[d+1]})/(d+1) \geq \text{rk}(\mathcal{B}^{[d]})/d$. See Lemma 5.7, Corollary 5.8, and Remark 4.2.

Question: compute $\lim_{d \rightarrow \infty} \text{rk}(\mathcal{B}^{[d]})/d$.

Remark: That NCFullRank is equivalent to deciding whether $\text{rk}(\mathcal{B}^{[d]}) = nd$ for some d was shown by Hrubeš and Wigderson [43]. Our formulation here is a straightforward quantitative generalization of their statement. Hrubeš and Wigderson's motivation was to study non-commutative arithmetic formulas with divisions.

- ($\mathbb{F} = \mathbb{C}$) Given $B_1, \dots, B_m \in M(n, \mathbb{Q})$, construct a completely positive operator $P : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$, sending $A \rightarrow \sum_{i \in [m]} B_i A B_i^\dagger$. For $c \in \mathbb{N}$, P is called rank c -decreasing, if there exists a positive semidefinite A , such that $\text{rk}(A) - \text{rk}(P(A)) = c$.

³As when the field size is large enough, the complement of the zero set of a nonzero polynomial is non-empty.

Question: compute the maximum c such that P is rank c -decreasing.

Remark: Gurvits stated the problem of deciding whether P is rank non-decreasing or not [39]. His original motivation was to study Edmonds' original (commutative) problem, and the main result in [39] solves the case when the commutative and non-commutative rank coincide (see Theorem 3.1).

4. (NCFullRank) Consider the action of $(A, C) \in \mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ on a tuple of matrices $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ by sending it to $(AB_1C^T, \dots, AB_mC^T)$.⁴ Let $R(n, m)$ be the \mathbb{F} -algebra of invariant polynomials with respect to this action. The *nullcone* of $R(n, m)$ is the common zero of all homogeneous positive-degree polynomials in $R(n, m)$.

Question: decide whether or not (B_1, \dots, B_m) is in the nullcone of $R(n, m)$.

That the original formulation is equivalent to (1) comes from [33]. The equivalence between (1) and (3) is straightforward. The equivalence among decision versions of (1) and (2), and (4) can be obtained via the ring of matrix semi-invariants, as described in Section 1.3. One way to prove the equivalence between (1) and (2) is via Theorem 5.11.

To summarize, the non-commutative Edmonds problem can be derived naturally from the perspectives of quantum information theory,⁵ and invariant theory. It is of great interest in non-commutative algebraic computation with divisions, and in the study of matrix spaces of low rank. Our motivation to study this is because a solution to the non-commutative Edmonds problem will throw light on its commutative counterpart. Shrunk subspaces form a natural and important witness for the singularity of a matrix space. Therefore, if the non-commutative Edmonds problem can be solved deterministically in polynomial time, it means that, for SDIT, the bottleneck lies in recognizing those singular matrix spaces without such witnesses. This connection will be detailed in Section 3.

1.3 Matrix semi-invariants

Formulation (1), (2) and (4) have a common origin, namely the invariant ring $R(n, m)$ described in (4). We shall call $R(n, m)$ the *ring of matrix semi-invariants*, as (1) it is closely related to the classical ring of matrix invariants [63] (see below for the definition, and [1, 23] for the precise relationship between these two rings); and (2) it is the ring of semi-invariants of the representation of the m -Kronecker quiver with dimension vector (n, n) . Here, the m -Kronecker quiver is the quiver with two vertices s and t , and m arrows pointing from s to t . When $m = 2$, it is the classical Kronecker quiver. The reader is referred to [21, 27, 66] for a description of the semi-invariants for arbitrary quivers.

The equivalence between (1) and (4) comes from the observation that the (B_1, \dots, B_m) with a shrunk subspace are exactly the points in the nullcone of $R(n, m)$ [1, 7]. The equivalence between (2) and (4) can be seen from the first fundamental theorem (FFT) of matrix semi-invariants [1, 21, 27, 66]. To describe this we need some notations: for $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. Note that $R(n, m) \subseteq \mathbb{F}[x_{i,j}^{(k)}]$ where $i, j \in [n]$, $k \in [m]$, and $x_{i,j}^{(k)}$ are independent variables. Let $X_k = (x_{i,j}^{(k)})_{i,j \in [n]}$ be a matrix of variables. Then for $A_1, \dots, A_m \in M(d, \mathbb{F})$, $\det(A_1 \otimes X_1 + \dots + A_m \otimes X_m)$ is a matrix semi-invariant, and every matrix semi-invariant is a linear combination of such polynomials. Therefore, (B_1, \dots, B_m) is in the nullcone, if and only if for all $d \in \mathbb{Z}^+$ and all $(A_1, \dots, A_m) \in M(d, \mathbb{F})^{\oplus m}$, $A_1 \otimes B_1 + \dots + A_m \otimes B_m$ is singular.

It is well-known that the matrix semi-invariant ring is finitely generated, by Hilbert's celebrated work [42]. This implies that there exists some integer d such that those matrix semi-invariants of degree no more than d define $R(n, m)$. This motivates the following definition.

Definition 1.1. $\beta(R(n, m))$ is the smallest integer d such that $R(n, m)$ is generated by invariants of degree $\leq d$.

⁴This action can also be written as: (A, C) sending (B_1, \dots, B_m) to $(AB_1C^{-1}, \dots, AB_mC^{-1})$. We adopt the transpose rather than the inverse, as the transpose yields a polynomial representation rather than a rational representation. Furthermore when Derksen's result is applied to the transpose, it gives a somewhat better bound (Fact 1.2).

⁵It remains to investigate the physical meaning for a super-operator to be rank non-decreasing though.

An explicit upper bound on $\beta(R(n, m))$ turns out to be particularly interesting for the purpose of the NCFullRank problem. As already suggested by Hrubes and Wigderson [43], if $R(n, m)$ has a degree bound $\beta = \beta(R(n, m))$, one can do the following: take m $d \times d$ variable matrices Y_1, \dots, Y_m , $Y_k = (y_{i,j}^{(k)})$ and form the polynomial

$$\det(Y_1 \otimes B_1 + \dots + Y_m \otimes B_m) \in \mathbb{F}[y_{i,j}^{(k)}]_{k \in [m], i,j \in [d]}. \quad (1.1)$$

Letting d go from 1 to β , this system of polynomials characterizes $\text{ncrk}(T) < n$: $\text{ncrk}(T) < n$ if and only if all these polynomials are the zero polynomial. This immediately gives a randomized algorithm for NCFullRank over large enough fields, with time complexity $\text{poly}(n, \beta)$.

In fact, for the above application, what really matters is another important bound $\sigma = \sigma(R(n, m))$. This is defined as the minimum integer d with the property that $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ is in the nullcone if and only if all polynomials of degree $\leq d$ in $R(n, m)$ vanish on $\{B_1, \dots, B_m\}$. It is clear that $\sigma \leq \beta$, and the above reasoning goes through when β is replaced by σ .

Over algebraically closed fields of characteristic 0, by directly employing Derksen's bounds for invariant rings satisfying certain general conditions [19], the following bound can be derived. For completeness we include a proof in Appendix A.

Fact 1.2 ([19]). *Over algebraically closed fields of characteristic 0, for $R(n, m)$, $\beta \leq \max\{2, 3/8 \cdot n^4 \cdot \sigma^2\}$, and $\sigma \leq 1/4 \cdot n^2 \cdot 4^{n^2}$.*

In particular, if σ is polynomial in n , then β is polynomial in n as well.

It is generally believed that over fields of characteristic 0, the bounds we get for $R(n, m)$ using Derksen's theorem is far from optimal. One reason to believe so is that $R(n, m)$ is closely related to another ring of invariants: let $A \in \text{SL}(n, \mathbb{F})$ act on $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ by simultaneous conjugation – i.e. A sends the tuple to $(AB_1A^{-1}, \dots, AB_mA^{-1})$. Denoted by $S(n, m)$, this is just the classical ring of matrix invariants [63]. The structure of $S(n, m)$ is well-understood. Over fields of characteristic 0, the first and second fundamental theorems for $S(n, m)$, and an n^2 upper bound for $\beta(S(n, m))$ were established in 1970's, by the works of Procesi, Razmysolov, and Formanek [32, 63, 64]. See [1, 23] for the precise relationship between the rings $R(n, m)$ and $S(n, m)$. Note that when applied to $S(n, m)$ over characteristic 0, Derksen's bound yields $\beta(S(n, m)) \leq \max\{2, 3/8 \cdot n^2 \cdot \sigma^2\}$ and $\sigma(S(n, m)) = n^{O(n^2)}$, far from the n^2 bound given above.

Another reason to believe that Derksen's bounds are far from optimal is that for certain small m or n , explicit generating sets of $R(n, m)$ have been computed in e.g. [22, 23, 25, 47]. In these cases, elements of degree $\leq n^2$ generate the ring.⁶

If we turn to positive characteristic fields then, to our best knowledge, no explicit bounds for $\beta(R(n, m))$ nor $\sigma(R(n, m))$ have been derived. Note here that the relation between β and σ as in Fact 1.2 is not known to hold, due to the assumption on the field properties there. This case is important, for example, in the application to identity testing, and division elimination for non-commutative arithmetic formulas with divisions over fields of positive characteristics [43]. For $S(n, m)$ over fields of positive characteristics, the FFT was established by Donkin in [28, 29]. Over fields of positive characteristic an $O(n^3)$ upper bound for $\sigma(S(n, m))$ can be derived from [12, Proposition 9], and in [24, 26] Domokos proved an upper bound $O(n^7 m^n)$ on $\beta(S(n, m))$.

1.4 Our results

In the previous sections, we defined the non-commutative Edmonds problem and the NCFullRank problem, and illustrated their connections to matrix semi-invariants. Indeed, our results suggest that progress on one topic helps to advance the other as well.

The first result shows that an upper bound for $\sigma(R(n, m))$ actually implies a *deterministic* algorithm for NCFullRank over \mathbb{Q} , rather than just a randomized one as in Section 1.3.

⁶We thank M. Domokos for pointing out this fact to us.

Proposition 1.3. *Over \mathbb{Q} , if the nullcone of $R(n, m)$ is defined by elements of degree $\leq \sigma = \sigma(n, m)$, then there exists a deterministic algorithm that solves NCFullRank with bit complexity polynomial in σ and the input size.*

In particular, if σ is a polynomial in n and m , then NCFullRank can be solved deterministically in polynomial time over \mathbb{Q} . The key ingredient here is Gurvits' algorithm for the Edmonds' problem, although that algorithm works only under a promise [39]. The distinction between deterministic and probabilistic is important: as illustrated at the end of Section 1.2, our original motivation of studying NCFullRank is to gain an understanding of SDIT, for which the question is to devise deterministic efficient algorithms.

Our main result is an algorithm that solves the non-commutative Edmonds problem using formulation (2). To ease the presentation, we give an informal statement of the main theorem in Section 5 (Theorem 5.11) here, and discuss its two consequences.

Theorem 1.4 (Theorem 5.11, informal). *Given a matrix space $\mathcal{B} \leq M(n, \mathbb{F})$ over a large enough field, there exists a deterministic algorithm that computes $\text{rk}(\mathcal{B})$ using $\text{poly}((n+1)!)$ many arithmetic operations. Over \mathbb{Q} the algorithm runs in time polynomial in the bit size of the input and $(n+1)!$.⁷*

In [33], Fortin and Reutenauer asked for “an algorithm which uses only linear-algebraic techniques.” Indeed, the algorithm for Theorem 1.4 may be viewed as one, though it relies on certain routines dealing with objects from cyclic field extensions and division algebras.

Two interesting consequences now follow. Firstly, we have a randomized efficient algorithm to compute the non-commutative rank if it differs from the commutative rank by a constant. (Recall that $\text{rk}(\mathcal{B}) \leq \text{ncrk}(\mathcal{B}) \leq 2\text{rk}(\mathcal{B})$.) Its easy proof is put after the statement of Theorem 5.11.

Corollary 1.5. *For $\mathcal{B} \leq M(n, \mathbb{F})$, let $c = \text{ncrk}(\mathcal{B}) - \text{rk}(\mathcal{B})$, and assume \mathbb{F} is of size $\Omega(n \cdot (n+1)!)$. Then the non-commutative rank of \mathcal{B} can be computed probabilistically in time polynomial in $(n+1)^{c+1}$.*

Secondly, we immediately obtain an explicit bound for $\sigma(R(n, m))$ as a consequence of Theorem 5.11. By Fact 1.2, we also get a bound on $\beta(R(n, m))$, over an algebraically closed field of characteristic 0. Its proof is also put after Theorem 5.11.

Corollary 1.6. *Over any field \mathbb{F} of size $\Omega(n \cdot (n+1)!)$, $\sigma(R(n, m)) \leq (n+1)!$. If furthermore \mathbb{F} is of characteristic 0 and algebraically closed, then $\beta(R(n, m)) \leq \max\{2, 3/8 \cdot n^4 \cdot ((n+1)!)^2\}$.*

This improves the bounds in Fact 1.2 over algebraically closed fields of characteristic 0. More importantly, to the best of our knowledge, this provides an explicit bound for $\sigma(R(n, m))$ over fields of positive characteristic for the first time. Furthermore to get this bound we only assume the field size to be large enough, whereas Fact 1.2 requires our field to be algebraically closed.

While the improvement from $2^{O(n^2)}$ to $2^{O(n \log n)}$ is modest, we believe it is nonetheless an interesting improvement from the technical point of view: note that the dimension of $\text{SL}(n, \mathbb{F})$ is $n^2 - 1$. In the line of research for bounds of an invariant ring R with respect to a group G (cf. [19, 62]), the dimension of G has to stand on the exponent for $\sigma(R)$, and to get a bound as $2^{o(\dim(G))}$ seems difficult there. Furthermore, the idea of using correctness of algorithms to get bounds on quantities of interest in invariant theory seems new, and may deserve to be explored further.

We also obtain certain structural results for $R(n, m)$, which are reported in [47].

1.5 More previous works

Connections between invariant theory and complexity theory. The results in this paper suggest a new link between invariant theory and complexity theory. Connections between the two fields have been emerging in recent years. We have already alluded to the direct connection with non-commutative arithmetic circuits, in

⁷All algorithms presented in this paper, when working over \mathbb{Q} , have bit complexity polynomial in the input size, and some additional parameters. Sometimes, we may omit the input size but focus on those more important parameters.

the work of Hrubeš and Wigderson [43] above. In a series of papers titled geometric complexity theory (GCT) [59, 60] (see also [8, 57]), Muliuley and Sohoni pointed out possible deep connections between problems in invariant theory and complexity theory. GCT addresses the fundamental lower bound problems in complexity theory, e.g. the permanent versus determinant problem, by linking them to problems in representation theory and algebraic geometry. In particular, in [58], Muliuley established a tight connection between derandomizing the Noether normalization lemma, and black-box derandomizing the polynomial identity test. The degree bounds of various invariant rings are of central importance in that work. We briefly remark that a polynomial bound for $\beta(R(n, m))$, if proven, will yield similar results as what the n^2 degree bound for $S(n, m)$ has yielded in [58].

More previous works on Edmonds' problem. Some earlier work on this problem was cited at the beginning of this article. Here we mention more related work.

An interesting instance of Edmonds' problem is the module isomorphism problem. Specifically, assume that we are given two n -dimensional modules U and U' for the free algebra \mathcal{A} over \mathbb{F} with k generators as k -tuples G_1, \dots, G_k and G'_1, \dots, G'_k of n by n matrices. Then $\text{Hom}_{\mathcal{A}}(U, U')$ is the \mathbb{F} -linear subspace of $\text{Hom}_{\mathbb{F}}(U, U')$, identified with $M(n, \mathbb{F})$, consisting of matrices X with $XG_i = G'_iX$ ($i = 1, \dots, k$). As these conditions are linear in the entries of X , the space $\text{Hom}_{\mathcal{A}}(U, U')$ can be obtained by solving a system of homogeneous linear equations in n^2 elements. Furthermore, U' is isomorphic to U if and only if there exists a nonsingular matrix in $\text{Hom}_{\mathcal{A}}(U, U')$. In turn, any such nonsingular matrix witnesses an isomorphism and, by the Schwartz-Zippel lemma, for sufficiently large base field a random homomorphism will be an isomorphism. Due to the special algebraic structure behind this problem, it can be solved even by deterministic polynomial-time methods, see the method of Chistov, Ivanyos and Karpinski [11] working over many fields, or a different approach of Brooksbank and Luks [6] which works over arbitrary fields, and an extension of the first method to arbitrary fields given by Ivanyos, Karpinski and Saxena in [45]. Interestingly, the general case of finding a surjective or injective homomorphism between non-isomorphic modules deterministically turns out to be as hard as the constructive version of Edmonds' general problem [45].

Recall that one motivation to study Edmonds' problem is due to its implications to certain combinatorial problems. This line of research mostly focuses in the case when the given matrices are of particular form, e.g. rank-1 and certain generalizations [36, 41, 45, 61] as used in bipartite graph matchings, or skew symmetric rank-2 and certain generalizations [35, 37, 38] as used in general graph matchings.

Another line of research deals with matrix spaces that satisfy certain properties. Note that properties of matrix spaces should not depend on a particular basis. For example, we can define a property of matrix spaces as, “having a basis consisting of rank-1 matrices.” So if \mathcal{B} has a basis consisting of rank-1 matrices, \mathcal{B} may not necessarily be presented using this rank-1 basis. We are not aware of any result on the complexity of finding rank-1 generators for rank-1 spanned matrix spaces, if it is given by a basis consisting of not necessarily rank-1 matrices. We believe that the problem is hard. Thus the results in [36, 41, 45, 61], which assume that the input is given by a rank-1 basis, do not translate to algorithms for rank-1 spanned matrix spaces.

As far as we are aware, there are two references for SDIT which assume only properties of matrix spaces. The first one is Gurvits' algorithm in [39]; this algorithm works over \mathbb{Q} , and assumes a property which Gurvits called “Edmonds-Rado.” His algorithm, put in the context of this paper, is rephrased as Theorem 3.1. Gurvits left open the problem of developing a deterministic efficient algorithm for rank-1 spanned matrix spaces over finite fields. This was settled in affirmative in [44], the other reference that assumes properties of matrix spaces.

Recall that the other major incentive to study Edmonds' problem is to understand arithmetic circuit lower bounds via [10, 48]. We believe that for this goal, a better indication of progress is to use properties of matrix spaces, rather than properties of the given matrices. One reason is that, whether a matrix space contains a nonsingular matrix, is a property of matrix spaces. Another reason is that many properties of matrix spaces seem difficult to test algorithmically. Furthermore, note that in this paper we heavily rely on algorithmic techniques developed in [39] and [44]. This may be viewed as another evidence of the importance of working with properties of matrix spaces.

Connections to Kronecker coefficients. Recently, there was an interest in studying the semi-invariants of the m -Kronecker quivers due to its connection with the Kronecker coefficients [1, 2, 55], namely the multiplicities in the direct sum decompositions of the tensor products of two irreducible representations of symmetric groups. Giving a positive combinatorial description of these coefficients is considered to be one of the most important problems in the combinatorial representation theory of symmetric groups.

1.6 Update on recent progress

There have been some exciting developments since we posted a version of this paper on the arXiv.

First, Garg et al. presented a deterministic polynomial-time algorithm for computing the non-commutative rank over \mathbb{Q} [34]. This is achieved via a closer analysis of Gurvits' algorithm [39]. Their analysis uses the exponential bounds on $\sigma(R(n, m))$ as proved in this paper, or deducible from Derksen's result [19]. It should be noted that the given algorithm is not constructive, in that it fails to produce a witness (e.g. shrunk subspaces).

Second, Derksen and Makam proved that $\sigma \leq n^2 - n$ over large enough fields [20]. Over fields of characteristic zero this implies $\beta(R(n, m)) = O(n^6)$, settling the question of whether there is a polynomial degree bound on the generators for this ring of invariants. To prove the upper bound on $\sigma(R(n, m))$, Derksen and Makam discover a concavity property of blow-ups, and rely crucially on Lemma 5.6 proved in this paper.

After [20] appeared, in [46] we show that the technique of Derksen and Makam can be constructivized. By combining that with the techniques in this paper, we obtain a constructive deterministic polynomial-time algorithm for computing the non-commutative rank over large enough fields. There we also present another independent proof of $\sigma(R(n, m)) \leq n^2 + n$. That argument also builds on Lemma 5.6 from this paper and is much simpler than the concavity argument of Derksen and Makam.

Organization. In Section 2 we present certain preliminaries. In Section 3 we give an exposition of the natural connection between commutative and the non-commutative Edmonds problem, and prove Proposition 1.3. In Section 4 we present an efficient construction of division algebras, to be used in proving the main technical lemma Lemma 5.4. In Section 5 we prove the formal version of Theorem 1.4 (Theorem 5.11) and deduce Corollary 1.5 and 1.6.

2 Preliminaries

2.1 Notation

For the reader's convenience, we collect the main notations in this section. Some of these were already introduced in the introduction.

For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. Given two vector spaces U and V , $U \leq V$ denotes that U is a subspace of V . $\mathbf{0}$ denotes the zero vector or the trivial vector space.

Let \mathbb{F} be a field. $\text{char}(\mathbb{F})$ denotes the characteristic of \mathbb{F} . $M(n, \mathbb{F})$ is the linear space of $n \times n$ matrices over \mathbb{F} . The rank of $A \in M(n, \mathbb{F})$ is denoted $\text{rk}(A)$. The corank of A , $\text{cork}(A)$, is $n - \text{rk}(A)$. For $U \subseteq \mathbb{F}^n$, $A^{-1}(U) = \{v \in \mathbb{F}^n \mid A(v) \in U\}$. I denotes the identity matrix.

A linear subspace of $M(n, \mathbb{F})$ is called a *matrix space*. For $B_i \in M(n, \mathbb{F})$, $i \in [m]$, $\langle B_1, \dots, B_m \rangle$ denotes the matrix space spanned by the B_i 's. For a matrix space $\mathcal{B} \leq M(n, \mathbb{F})$, $\text{rk}(\mathcal{B})$ is defined as $\max\{\text{rk}(B) \mid B \in \mathcal{B}\}$. We call \mathcal{B} *singular*, if $\text{rk}(\mathcal{B}) < n$. For $\mathcal{B} \leq M(n, \mathbb{F})$ and $U \leq \mathbb{F}^n$, $\mathcal{B}(U) := \langle \cup_{B \in \mathcal{B}} B(U) \rangle$. The non-commutative rank, $\text{ncrk}(\mathcal{B})$, is $n - c$ where c is the maximum integer such that there exists a c -shrunk subspace.

For $A \in M(d, \mathbb{F})$ and $B \in M(n, \mathbb{F})$, the tensor product $A \otimes B$ is a $d \times d$ block matrix with block size $n \times n$. For $A = (a_{i,j})_{i,j \in [d]}$, the (i, j) -th block of $A \otimes B$ is $a_{i,j}B$. For $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, the d th tensor blow-up, $\mathcal{B}^{[d]} := M(d, \mathbb{F}) \otimes \mathcal{B} \leq M(dn, \mathbb{F})$. A linear basis of $\mathcal{B}^{[d]}$ is $\{E_{i,j} \otimes B_k \mid i, j \in [d], k \in [m]\}$, where $E_{i,j}$ is the matrix with 1 at the (i, j) th position, and 0 otherwise. In Section 5 it will be easier to work with $\mathcal{B}^{\{d\}} := \mathcal{B} \otimes M(d, \mathbb{F})$. As $\mathcal{B}^{[d]} \cong \mathcal{B}^{\{d\}}$, the latter will also be referred to as the d th tensor blow-up of \mathcal{B} .

2.2 The second Wong sequences

Let us introduce a key tool to be used in Section 5, called the second⁸ (generalized) Wong sequence. This was used by Fortin and Reutenauer [33], and rediscovered by the first two authors with Karpinski and Santha in [44] to solve Edmonds' problem for rank-1 spanned matrix spaces over arbitrary fields.

Given $A \in M(n, \mathbb{F})$ and $\mathcal{B} \leq M(n, \mathbb{F})$, the second Wong sequence of (A, \mathcal{B}) is the following sequence of subspaces in \mathbb{F}^n : $W_0 = \mathbf{0}$, $W_1 = \mathcal{B}(A^{-1}(W_0))$, \dots , $W_i = \mathcal{B}(A^{-1}(W_{i-1}))$, \dots . It can be proved that $W_0 < W_1 < W_2 < \dots < W_\ell = W_{\ell+1} = \dots$ for some $\ell \in \{0, 1, \dots, n\}$. W_ℓ is then called the limit of this sequence, denoted as W^* .

A useful way to understand the second Wong sequence is to view it as a linear algebraic analogue of the augmenting path on bipartite graphs. While not precise, we find this intuition helpful. That is, we view matrices as linear maps from V to W , $V \cong W \cong \mathbb{F}^n$. Vectors in V and W may be thought of as the “vertices” on the left and right part, respectively. Then for $A \in \mathcal{B}$, thinking of A as a given matching, $A^{-1}(\mathbf{0})$ can be understood as identifying those “vertices unmatched by A on the left part.” Then $\mathcal{B}(A^{-1}(\mathbf{0}))$ is understood as taking those “edges” outside A , and $A^{-1}(\mathcal{B}(A^{-1}(\mathbf{0})))$ is understood as taking a further step with those “edges” in A . And so on.

The key fact is that, when $A \in \mathcal{B}$, $W^* \leq \text{im}(A)$ if and only if there exists a $\text{cork}(A)$ -shrunk subspace [44, Lemma 9] (reproduced below as Fact 2.1). If this is the case, A is of maximum rank and $A^{-1}(W^*)$ is a $\text{cork}(A)$ -shrunk subspace. It is clear that the second Wong sequence can be computed using polynomially many arithmetic operations. The direct way to compute the second Wong sequences over \mathbb{Q} may cause the bit lengths to explode. If testing whether $W^* \leq \text{im}(A)$ is the only concern (as in our application here), by replacing A^{-1} with some appropriate “pseudo-inverse” of A , the bit lengths of the intermediate numbers up to the first W_k , $W_k \not\leq \text{im}(A)$, can be bounded by a polynomial of the input size. We refer the reader to [44, Lemma 10] for this trick.

When $|\mathbb{F}|$ is $\Omega(n)$, this immediately gives a method to decide whether $\text{ncrk}(\mathcal{B}) = \text{rk}(\mathcal{B})$ as in [33]: randomly choose a matrix $A \in \mathcal{B}$, which will be of maximal rank with high probability. Then compute the second Wong sequence of (A, \mathcal{B}) and check whether the limit $W^* \subseteq \text{im}(A)$.

For completeness we summarize the above discussion together as a fact.

Fact 2.1 ([44, Lemmas 9 and 10]). *Let $A \in \mathcal{B} \leq M(n, \mathbb{F})$, and let W^* be the limit of the second Wong sequence of (A, \mathcal{B}) . Then there exists a $\text{cork}(A)$ -shrunk subspace of \mathcal{B} if and only if $W^* \subseteq \text{im}(A)$. If this is the case then $A^{-1}(W^*)$ is a $\text{cork}(A)$ -shrunk subspace of \mathcal{B} . In the algebraic RAM model as well as over \mathbb{Q} we can detect whether $W^* \subseteq \text{im}(A)$ and if so compute a shrunk subspace in deterministic polynomial time.*

For a matrix space \mathcal{B} of dimension 2, $\text{rk}(\mathcal{B}) = \text{ncrk}(\mathcal{B})$ for large enough \mathbb{F} ; this follows from the Kronecker-Weierstrass theory of matrix pencils, and alternative proofs may be found in [4, 31]. Due to this fact, it was observed in [44] that by utilizing the second Wong sequence we have the following.

Fact 2.2 ([44, Fact 11]). *Assume that $|\mathbb{F}| > n$, and let $\mathcal{B} = \langle A, B \rangle \leq M(n, \mathbb{F})$. Then $\text{rk}(A) = \text{rk}(\mathcal{B})$ if and only if for any $i \in [n]$, $(\mathcal{B}A^{-1})^i(\mathbf{0}) \leq \text{im}(A)$.*

3 Gurvits' algorithm and Proposition 1.3

Commutative and the non-commutative Edmonds problem: a natural pair. Viewing matrices as linear maps between two vector spaces, one may suspect Edmonds problem to be a linear algebraic analogue of the maximum matching problem on bipartite graphs, with elements of the underlying vector spaces as being the left and right side vertices, and the matrices as giving us edges – mapping a vector on the left side to one on the right side. Given such a correspondence, one may ask whether an analogue of Hall's theorem holds in this setting, i.e., is it true that a matrix space either has a matrix of rank s , or has an $(n-s+1)$ -shrunk subspace;

⁸The first Wong sequence is the dual of the second one; this naming convention is due to Wong who in [69] defined the two sequences for the special case $\mathcal{B} = \langle B \rangle$.

or put differently, whether $\text{rk}(\mathcal{B}) = \text{ncrk}(\mathcal{B})$ holds for all \mathcal{B} . This is far from the truth! For example, for the space of skew-symmetric matrices ($A = -A^T$) of size 3, we have $\text{ncrk} = 3$ and $\text{rk} = 2$.

That is, while for the bipartite maximum matching problem, matchings and shrunk subsets are two sides of the same coin, in the linear algebraic setting, this coin splits into two problems: Edmonds' original (commutative) problem asks to compute the maximum rank, and the non-commutative Edmonds problem asks to compute the maximum c for the existence of a c -shrunk subspace.

Rank-1 spanned matrix spaces. Now we point out that several results on the (commutative) Edmonds problem can be viewed, and should be understood as, resolving the non-commutative counterpart. For this, note that shrunk subspaces are a natural witness for the singularity of matrix spaces: this construction can be dated back to 1930's in T. G. Room's book [65], and plays a key role in several results which solve special cases of Edmonds' problem including [39, 44, 52].

A particular case of interest is rank-1 spanned matrix spaces: those matrix spaces that have a basis consisting of rank-1 matrices. For rank-1 spanned spaces, the analogue of Hall's theorem holds [52], so the commutative and the non-commutative Edmonds problems coincide. Therefore, the known results for rank-1 spanned spaces [39, 44] can be viewed as solving either the non-commutative Edmonds problem or the commutative one. In retrospect, the results for rank-1 spanned spaces rely on shrunk subspaces in such a critical way that they should be understood as solving NCFullRank rather than the commutative version for this special case:

- The core of Gurvits' algorithm [39] is an iterative procedure called the operator Sinkhorn's scaling procedure. When applied to a matrix space \mathcal{B} , this procedure converges, if and only if \mathcal{B} has a shrunk subspace.
- The key tool in [44] is the second Wong sequence as described in Fact 2.1. When applied to $A \in \mathcal{B} \leq M(n, \mathbb{F})$, this sequence stabilizes in polynomially many number of steps, and the limit subspace is contained in $\text{im}(A)$ if and only if \mathcal{B} has an $\text{cork}(A)$ -shrunk subspace.

Gurvits' algorithm; Proof of Proposition 1.3. In fact, Gurvits' algorithm works by assuming that an analogue of Hall's theorem for perfect matchings holds.

Theorem 3.1 ([39]). *Over \mathbb{Q} , given a matrix space \mathcal{B} such that either $\text{rk}(\mathcal{B}) = n$ or $\text{ncrk}(\mathcal{B}) < n$, there exists a deterministic polynomial-time algorithm that solves SDIT, and therefore NCFullRank.*

Gurvits' algorithm almost solves NCFullRank over \mathbb{Q} . The only problem is that for a matrix space \mathcal{B} with $n = \text{ncrk}(\mathcal{B}) > \text{rk}(\mathcal{B})$ the algorithm may give a wrong answer. (We note that even when the input to the algorithm is such a matrix space it terminates in polynomially many steps.) We observe that this can be rectified by considering matrix semi-invariants up to the upper bound for $\sigma(R(n, m))$.

Proof of Proposition 1.3. Recall that, by assumption, the nullcone of $R(n, m)$ is defined by elements of degree $\leq \sigma(R(n, m))$ over \mathbb{Q} . Also, given a matrix space $\mathcal{B} \in M(s, \mathbb{Q})$, Gurvits' algorithm either reports that $\text{rk}(\mathcal{B}) = s$, or $\text{ncrk}(\mathcal{B}) < s$. When $\text{rk}(\mathcal{B}) = s$ or $\text{ncrk}(\mathcal{B}) < s$, it is always correct.

The algorithm is easy to describe: for $d = 1, \dots, \sigma = \sigma(R(n, m))$ run Gurvits' algorithm with input $\mathcal{B}^{[d]}$. If for some d , Gurvits' algorithm reports $\text{rk}(\mathcal{B}^{[d]}) = dn$, then output $\text{ncrk}(\mathcal{B}) = n$ and halt. Otherwise, return $\text{ncrk}(\mathcal{B}) < n$.

It is clear that this algorithm runs in time polynomial in the input size and σ . Note that a linear basis of $\mathcal{B}^{[d]}$ can be constructed easily in time polynomial in the input size of \mathcal{B} and d .

From the discussion in Section 1.3, the correctness is also easy to see. Specifically, assuming the bound on σ , the simultaneous vanishing of $\det(Y_1 \otimes B_1 + \dots + Y_m \otimes B_m) \in \mathbb{Q}[y_{i,j}^{(k)}]$ for $d = 1, \dots, \sigma$, characterizes whether $\text{ncrk}(\mathcal{B}) < n$ or not. Therefore, if $\text{ncrk}(\mathcal{B}) = n$, then for some $d \leq \sigma$, $\text{rk}(\mathcal{B}^{[d]})$ is full. On the other hand if $\text{ncrk}(\mathcal{B}) < n$, then there is a shrunk subspace U . For each d , $\mathbb{Q}^d \otimes U$ is a shrunk subspace of $\mathcal{B}^{[d]}$ and so $\text{ncrk}(\mathcal{B}^{[d]}) < dn$ for any d . \square

Implications to Gurvits' algorithm. The invariant-theoretic viewpoint also connects to a question of Gurvits in [39]. In [39], given a basis $\{B_1, \dots, B_m\}$ of $\mathcal{B} \leq M(n, \mathbb{C})$, Gurvits associates with it a completely positive operator i.e., a linear map $F : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$. The main algorithmic technique is the so-called operator Sinkhorn's iterative scaling procedure, which is applied to F . This procedure is a quantum generalization of the classical Sinkhorn's iterative scaling procedure, which is applied to nonnegative matrices, and can be used to approximate the permanent, and to decide the existence of perfect matchings [40, 51]. Gurvits proved that this procedure, when applied to the operator T derived from a matrix space \mathcal{B} , converges if and only if \mathcal{B} has a shrunk subspace. He proved this using a continuous but non-differentiable function, called the capacity of an operator, denoted as $\text{Cap}(F)$. Specifically, he showed that \mathcal{B} has a shrunk subspace if and only if $\text{Cap}(F) = 0$. Gurvits asked whether there exists a “nice” function, like a polynomial with integer coefficients, that characterizes \mathcal{B} with shrunk subspaces. Our previous argument suggests that there exists a set of polynomial functions with integer coefficients, whose simultaneous vanishing characterizes those \mathcal{B} with shrunk subspaces, and therefore a “nice” substitute for Gurvits' capacity. However, the number of these polynomial functions depends on the degree bound for matrix semi-invariants.

4 Efficient construction of division algebras

Division algebras, and efficient construction of such algebras with explicit matrix representations play a crucial role in the main technical lemma, Lemma 5.7, in this paper. In this section we present an efficient construction of such algebras based on Kummer extensions.

4.1 Basic facts about central division algebras

Let us first introduce some basic facts about central division algebras. Proofs of these statements can be found in [50].

Let \mathbb{F} be a field. A *division algebra* D over \mathbb{F} is an associative \mathbb{F} -algebra in which the non-zero elements are invertible. The *center* of a division algebra D over \mathbb{F} is obviously an extension field of \mathbb{F} . All the division algebras considered in this section are finite dimensional over their center. The *opposite division algebra* D^{op} is the algebra with the same set of elements as D and with multiplication $x \cdot y$ defined to be $y * x$, with $*$ being the multiplication in D .

When the center coincides with \mathbb{F} , we say that D is a *central division algebra* over \mathbb{F} , and in this case, $\dim_{\mathbb{F}}(D) = d^2$ for some positive integer d . This d is called the *index* of D . For $x, y \in D$ we can consider the linear transformation $\mu_{x,y}$ on D , considered as a vector space of dimension d^2 over \mathbb{F} , defined as $\mu_{x,y}z = x * z * y$. The linear extension of the map $x \otimes y \mapsto \mu_{x,y}$ to $D \otimes D^{op}$ gives an isomorphism $D \otimes D^{op} \cong M(d^2, \mathbb{F})$, the algebra of $d^2 \times d^2$ matrices with entries in \mathbb{F} . The image of $D \otimes I$ under this isomorphism gives a representation of D in the space of $d^2 \times d^2$ matrices with entries in \mathbb{F} . What is important for us is the observation that matrices giving us a representation of D^{op} in $M(d^2, \mathbb{F})$ commute with the matrices giving us a representation of D in $M(d^2, \mathbb{F})$.

4.2 Constructing cyclic field extensions under a coprime condition

Our division algebras will be cyclic algebras, that is, non-commutative algebras constructed from cyclic extensions of fields. In this subsection we present an efficient construction of such field extensions, under the condition that the extension degree and the field characteristic are coprime.

Recall that a *cyclic extension of a field* \mathbb{K} is a finite Galois extension of \mathbb{K} having a cyclic Galois group. By constructing a cyclic extension \mathbb{L} we mean constructing the extension as an algebra over \mathbb{K} , e.g., by giving an array of *structure constants* with respect to a \mathbb{K} -basis for \mathbb{L} defining the multiplication on \mathbb{L} as well as specifying a generator of the Galois group, e.g., by its matrix with respect to a \mathbb{K} -basis. Recall that for a finite dimensional algebra \mathcal{A} over the field \mathbb{K} , a common way to specify the multiplication is using an array of structure constants with respect to a \mathbb{K} -basis A_1, \dots, A_d . These are d^3 elements γ_{ijk} of \mathbb{K} such that $A_i A_j = \sum_{k=1}^d \gamma_{ijk} A_k$. Then we can represent elements of \mathcal{A} by the vectors of their coordinates in terms of

the basis A_1, \dots, A_d . The size of the data representing the structure constants gives some control over the size of the data representing the product of elements. For example, consider the following situation: \mathbb{K} is the function field $\mathbb{F}'(Z)$, where \mathbb{F}' is a field and Z a formal variable. The structure constants happen to be polynomials in $\mathbb{F}'[Z]$. Then for two elements of \mathcal{A} with their coordinates being polynomials in $\mathbb{F}'[Z]$, their product will have also polynomial coordinates, and the degrees of the coordinates of the product are upper bounded by the sum of the maximum degrees of coordinates of the factors, plus the maximum degree of the structure constants.

Lemma 4.1. *Let \mathbb{F}' be a field. Let d be any non-negative integer if the characteristic of \mathbb{F}' is zero, otherwise assume that d is not divisible by the characteristic of \mathbb{F}' . Assume that \mathbb{F}' contains a known primitive d th root of unity ζ , and let X be a formal variable. Then a cyclic extension \mathbb{L} having degree d over $\mathbb{K} := \mathbb{F}'(X)$ can be computed using $\text{poly}(d)$ arithmetic operations. \mathbb{L} will be given by structure constants with respect to a basis, and the matrix for a generator of the Galois group of \mathbb{L}/\mathbb{K} in terms of the same basis will also be given. All the output entries (the structure constants as well as the entries of the matrix representing the Galois group generator) will be polynomials of degree $\text{poly}(d)$ in $\mathbb{F}'[X]$. Furthermore for $\mathbb{F}' = \mathbb{Q}[\sqrt[d]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is $\text{poly}(d)$.*

Proof. Put $\mathbb{L} = \mathbb{F}'(Y_1)$ where $X = Y_1^d$. Then $1, Y_1, \dots, Y_1^{d-1}$ are a $\mathbb{F}'(X)$ -basis for \mathbb{L} with $Y_1^i Y_1^j = Y_1^{i+j}$ if $i + j \leq d$ and XY_1^{i+j-d} otherwise. Further note that the linear extension σ of the map sending Y_1^j to $\zeta^j Y_1^j$ is a $\mathbb{K}(X)$ -automorphism of degree d . \square

Remark 4.2. The construction above is known in the literature as a Kummer extension. When the characteristic is a prime p and a divisor of d , say $d = p^e d'$ where d' is prime to p , the Kummer extension \mathbb{K} should be replaced by a cyclic extension which is a product of a Kummer extension of degree d' and a cyclic extension of degree p^e described by Artin, Schreier and Witt [68]. Investigating the complexity of computing such extensions requires some further work. In [46], we conduct such a research and present an efficient construction of such extensions. The consequence on results in this paper will be reported in ibid..

4.3 Constructing cyclic division algebras

The following statement connects cyclic field extensions with central division algebras. It follows from Wedderburn's theorem characterizing cyclic division algebras (see e.g. [50, Theorem (14.9)]) as shown on Page 221 of [50].

Fact 4.3. *Let \mathbb{L} be a cyclic extension of degree d of a field \mathbb{K} . Let σ be a generator of the Galois group, and Z a number transcendental over \mathbb{L} . For the transcendental extension $\mathbb{L}(Z)$ of \mathbb{L} , σ extends to an automorphism (denoted again by σ) of $\mathbb{L}(Z)$ such that the fixed field of σ is $\mathbb{K}(Z)$. Thus $\mathbb{L}(Z)$ is a cyclic extension of $\mathbb{K}(Z)$. Consider the $\mathbb{K}(Z)$ -algebra D generated by (a basis for) \mathbb{L} and by an element U with relations $U^d = Z$ and $Ua = a^\sigma U$ (for every $a \in \mathbb{L}(Z)$, or, equivalently for every a from a fixed \mathbb{K} -basis for \mathbb{L}). Then D is a central division algebra of index d over $\mathbb{K}(Z)$.*

The following proposition is an algorithmic realization of Fact 4.3.

Proposition 4.4. *Let \mathbb{L} be a cyclic extension of degree d of a field \mathbb{K} , and suppose that \mathbb{L} is given by structure constants with respect to a \mathbb{K} -basis A_1, \dots, A_d . Similarly, a generator σ for the Galois group is assumed to be given by its matrix in terms of the same basis. Let Y be a formal variable. Then one can construct a $\mathbb{K}(Y)$ -basis Γ of $M(d, \mathbb{K}(Y))$ such that the $\mathbb{K}(Y^d)$ -linear span of Γ is a central division algebra over $\mathbb{K}(Y^d)$ of index d , using $\text{poly}(d)$ arithmetic operations in \mathbb{K} .*

Proof. Let $Z = Y^d$. Let D be a central division algebra over $\mathbb{K}(Z)$ as in Fact 4.3. The existence of a $\mathbb{K}(Z)$ -subalgebra D' of $M(d, \mathbb{K}(Y))$ isomorphic to D follows, e.g., from Theorem (14.7) of [50]. To construct a basis Γ for such a matrix algebra D' efficiently, note that $A_i U^j$, $i, j = 1, \dots, d$, form a $\mathbb{K}(Z)$ -basis of D . This is also a $\mathbb{K}(Y)$ -basis for the algebra $D'' = \mathbb{K}(Y) \otimes_{\mathbb{K}(Z)} D$. Consider also the element $U_0 = \frac{1}{Y} \otimes U$. Then $U_0^d = 1$. As the elements U_0^j are linearly independent over $\mathbb{K}(Y)$ and hence over $\mathbb{K}(Z)$ as well, we have that

$E = U_0 + U_0^2 + \dots + U_0^{d-1} + U_0^d$ is nonzero. As $U^j E = Y^j E$, we have $A_i E$ ($i = 1, \dots, d$) form a $\mathbb{K}(Y)$ -basis for the left ideal $D''E$ of dimension d . Now the action of D'' on this left ideal gives a matrix representation for D'' . Let γ_{kij} be the structure constants for the multiplication of \mathbb{L} (and of $\mathbb{L}(Z)$):

$$A_k A_i = \sum_{j=1}^d \gamma_{kij} A_j \quad (k, i = 1, \dots, d).$$

Also, let $\delta_{\ell ij}$ be the entries of the matrix of the ℓ th power of the generator σ of the Galois group:

$$A_i^{\sigma^\ell} = \sum_{j=1}^d \delta_{\ell ij} A_j \quad (\ell, i = 1, \dots, d).$$

(Notice that the matrix $(\delta_{\ell ij})_{ij}$ is the ℓ th power of $(\delta_{1ij})_{ij}$, whence the degrees of its elements are also bounded by $\text{poly}(d)$.) Then

$$A_k A_i E = \sum_{j=1}^d \gamma_{kij} A_j E$$

and

$$U^\ell A_i E = A_i^{\sigma^\ell} U^\ell E = A_i^{\sigma^\ell} Y^\ell E = Y^\ell \sum_{j=1}^d \delta_{\ell ij} A_j E.$$

Thus the matrix of the action of A_k has entries γ_{kij} and the matrix of the action of U^ℓ has entries $Y^\ell \delta_{\ell ij}$. Then the action of $U^\ell A^k$ can be obtained as the product of these two matrices. Let Γ consist of all such d^2 products, and the proof is concluded. \square

Combining Lemma 4.1 and Proposition 4.4, we immediately obtain the following.

Lemma 4.5. *Let $d, X, \mathbb{F}', \mathbb{K} = \mathbb{F}'(X)$ and \mathbb{L} be as in Lemma 4.1. In particular, if $\text{char}(\mathbb{K}) = p > 0$ then $p \nmid d$. Then one can construct a $\mathbb{F}'(X, Y)$ -basis Γ of $M(d, \mathbb{F}'(X, Y))$ such that the $\mathbb{F}'(X, Y^d)$ -linear span of Γ is a central division algebra over $\mathbb{F}'(X, Y^d)$ of index d , using $\text{poly}(d)$ arithmetic operations in \mathbb{F}' . In particular, for any $A \in \Gamma$, the entries of A are polynomials in $\mathbb{F}'[X, Y]$ of degree $\text{poly}(d)$. Furthermore for $\mathbb{F}' = \mathbb{Q}[\sqrt[d]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is also $\text{poly}(d)$.*

4.4 Some algorithmic issues for actual applications

To put the above construction in action, we need to handle a few algorithmic problems as follows.

4.4.1 Algorithmic issues when working with field extensions

Lemma 4.1 assumes the field \mathbb{F}' contains a known primitive d th root of unity ζ , where if $\text{char}(\mathbb{F}') = p > 0$ then $p \nmid d$. In actual applications, we may start with a field \mathbb{F} without a primitive d th root of unity in it, and attach one symbolically, which we still denote by ζ . However, this may cause some problem. Namely, constructing $\mathbb{F}' = \mathbb{F}[\zeta]$ would require factoring the polynomial $x^d - 1$ over \mathbb{F} , a task which cannot be accomplished using basic arithmetic operations. To see that this is indeed an issue notice that a black-box field may contain certain “hidden” parts of cyclotomic fields. Of course, over certain concrete fields, such as the rationals, number fields or finite fields of small characteristics, this can be done in polynomial time. However, even over finite fields of large characteristic no deterministic polynomial time solution to this task is known at present.

To get around this issue, one can perform the required computations over an appropriate factor algebra R of the algebra $C = \mathbb{F}[x]/(x^d - 1)$ in place \mathbb{F}' as if R were a field. To be specific, as d is not divisible by the characteristic, we know that C is semisimple – actually it is isomorphic to a direct sum of ideals, each of which is isomorphic to the splitting field $\mathbb{F}[\sqrt[e]{1}]$ of the polynomial $x^e - 1$ for some divisor e of d , and the

projection of x to such an ideal is a primitive e th root of unity. It follows that if we compute the ideal J generated by annihilators of $x^e - 1$, for all e a proper divisor of d , then $R = C/J$ is isomorphic to the direct sum of copies of the splitting field \mathbb{F}' of $x^d - 1$, and the projection of x to each component is a primitive d th root of unity. And this property is inherited by any proper factor of R . A computation using R instead of \mathbb{F}' may fail only at a point where we attempt to invert a non-invertible element of R . However, such an element must be a zero divisor. When this situation occurs, we replace R with the factor of R by its ideal generated by the zero divisor and restart the computation. Such a restart can clearly happen at most $d - 2$ times.

We explain what the above scheme entails in our actual tasks.

As the methods for Proposition 4.4 and Lemma 4.5 do not require division, the zero divisor issue does not occur there. Replacing \mathbb{F}' with R , the outcome Γ of Lemma 4.5 will actually be a free $R(X, Y)$ -basis for an algebra which is a direct sum of isomorphic copies of a division algebra, embedded into $M(d, R(X, Y))$.

Now consider the task of computing the rank of $M(N, \mathbb{F}')$. Note that we cannot talk about the “rank” of matrices in $M(N, R)$ which is not well-defined. But since R is a direct sum of \mathbb{F}' , the decomposition of R induces a decomposition of $M(N, R)$ into a direct sum of copies of $M(N, \mathbb{F}')$. We call the images of the projections of a matrix $B \in M(N, R)$ to the direct summands the *components* of B . The following lemma describes how to compute the maximum rank over the components.

Lemma 4.6. *Let R and \mathbb{F}' be as above, and suppose we are given a matrix $B \in M(N, R)$. Then there exists a deterministic polynomial-time algorithm that computes the maximum rank over the components of B .*

Proof. This can be achieved by combining division-free algorithms for computing the determinant by e.g. Kaltofen [49] (see also [53] for more such algorithms), and the parallel algorithm for computing the rank of a matrix by Mulymuley [56].

We include a sketch here for completeness. To start with, instead of B we consider the symmetric matrix $B' = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}$. Then let x and y be two formal variables. Form a matrix $D = \text{diag}(1, y, y^2, \dots, y^{2N-1})$, and compute $\det(xI - DB')$ using [49], considered as a polynomial in $\mathbb{F}'(y)[x]$. Let M be the maximum integer such that x^M divides $\det(xI - DB')$, and return $(2N - M)/2$.

By [56], the above procedure on a matrix from $M(N, \mathbb{F}')$ returns its rank. Now for $B \in M(N, R)$, since it is (implicitly) a direct sum of several copies of $M(N, \mathbb{F}')$, the above algorithm on B can be viewed as working with these components “in parallel”, and the resulting $\det(xI - DB)$ is a direct sum of $\det(xI - DB'_i)$ where B'_i are the components of B' . It is then not hard to deduce that the above procedure computes the maximum rank over the components of B . \square

Remark 4.7. Using the method of Lemma 4.6 for rank computations, we will obtain an algorithm that does not require division in R at all, and hence we will not need the above mentioned restarts. Another possibility would be doing Gaussian elimination and restarting the computation once a zero divisor is met as described. If no zero divisors are met and the rank is r , then it means that the columns of B generated a free module over R of rank r , so each component is also of rank r over \mathbb{F}' .

Finally, we note that a similar issue, namely that a black box field may even contain infinite algebraic extensions of its subfields has been circumvented by using the transcendental extension $\mathbb{K} = \mathbb{F}'(X)$ in the construction of cyclic extensions (Lemma 4.1).

4.4.2 Computing the rank of matrices over a rational function field in few variables

Note that the matrices from Lemma 4.5 are matrices over a rational function field. Therefore we will need to compute the rank of matrices in such form.

Proposition 4.8. *Let \mathbb{F}' be a field and $\mathbb{K} = \mathbb{F}'(X_1, X_2, \dots, X_k)$ be a pure transcendental extension of \mathbb{F}' . Let A be an $N \times N$ matrix with entries as quotients of polynomials from $\mathbb{F}'[X_1, X_2, \dots, X_k]$, where the polynomials are explicitly given as sums of monomials. Assume that the degrees of the polynomials appearing in A are upper bounded by D . If $|\mathbb{F}'| = (ND)^{\Omega(k)}$, then we can find in time $(ND)^{O(k)}$ a matrix $B \in M(N, \mathbb{F}')$ with $\text{rk}(B) = \text{rk}(A)$.*

In particular, if k is a constant – $k = 2$ as used in Lemma 5.3 for the procedure in Lemma 5.7 – then the above procedure runs in polynomial time.

Proof. We multiply the entries of A by an easily computable common multiple (e.g., the product) of their denominators to obtain a matrix with polynomial entries from $\mathbb{F}'[X_1, X_2, \dots, X_k]$. The data describing this matrix has size polynomial in the size of the input data. In particular, the degree of the determinant of any sub-matrix is upper bounded by a polynomial s in $(ND)^k$. We have assumed that $|\mathbb{F}'| = (ND)^{\Omega(k)}$. Then from $(ND)^{O(k)}$ specializations by elements of a subset of size $s + 1$ of \mathbb{F}' , at least one gives a matrix with entries \mathbb{F}' having the same rank as the original matrix. Thus the rank of A can be computed by computing the rank of $(ND)^{O(k)}$ matrices over \mathbb{F}' . \square

Note that if we simulate \mathbb{F}' using R as described in Section 4.4.1 then we shall apply the procedure in Lemma 4.6 after specializing the variables as in Proposition 4.8.

5 Finding a nonsingular matrix in blow-ups

In this section we describe, given $\mathcal{B} \leq M(n, \mathbb{F})$, how to compute a nonsingular matrix in $\mathcal{B}^{[d]} = M(d, \mathbb{F}) \otimes \mathcal{B}$ for some $d \leq (n + 1)!$, or certify that this is not possible.

One important note is due here: as per our notation, matrices in $M(d, \mathbb{F}) \otimes M(n, \mathbb{F})$ are viewed as block matrices, where each block is of size $n \times n$. This is more convenient when describing semi-invariants. In this section, it will be more convenient to work with $M(n, \mathbb{F}) \otimes M(d, \mathbb{F})$, namely each block is of size $d \times d$. This is consistent with other parts simply because $M(d, \mathbb{F}) \otimes M(n, \mathbb{F}) \cong M(n, \mathbb{F}) \otimes M(d, \mathbb{F})$. Therefore, we identify $M(nd, \mathbb{F}) \cong M(n, \mathbb{F}) \otimes M(d, \mathbb{F})$, and fix such a decomposition. Recall the notation $\mathcal{B}(U)$, and $\mathcal{B}^{\{d\}}$ from Section 2.1.

After some preparation, we prove the main technical lemma called the regularity lemma for blow-ups. Then we prove Theorem 5.11 (the formal version of Theorem 1.4), and Corollaries 1.5 and 1.6 follow easily.

5.1 Preparations

A characterization of blow-ups.

Proposition 5.1. *For $\mathcal{A} \leq M(dn, \mathbb{F})$, $\mathcal{A} = \mathcal{B}^{\{d\}}$ for some $\mathcal{B} \leq M(n, \mathbb{F})$ if and only if $(I \otimes M(d, \mathbb{F}))\mathcal{A}(I \otimes M(d, \mathbb{F})) = \mathcal{A}$.*

We remind the reader that $(I \otimes M(d, \mathbb{F}))\mathcal{A}(I \otimes M(d, \mathbb{F})) = \{XAY : X \in I \otimes M(d, \mathbb{F}), A \in \mathcal{A}, Y \in I \otimes M(d, \mathbb{F})\}$. Note that multiplication of tensor products of matrices obeys the rule $(X_1 \otimes Y_1)(X_2 \otimes Y_2) = X_1X_2 \otimes Y_1Y_2$.

Proof. The only if part is obvious. To see the reverse implication, for $i, j \in [d]$, let E_{ij} stand for the elementary matrix in $M(d, \mathbb{F})$ in which the (i, j) th entry is 1 and the others are zero. Then for every quadruple $(i, j, i', j') \in [d]^4$ we have $\sum_{k=1}^d E_{ki}E_{i'j'}E_{jk} = \delta_{ii'}\delta_{jj'}I$, where $\delta_{\ell\ell'}$ stands for the Kronecker delta. Any element A of \mathcal{A} can be written as $\sum_{i,j=1}^d B_{ij} \otimes E_{ij}$ with $B_{ij} \in M(n, \mathbb{F})$. Then for every $i, j \in [d]$ we have $\sum_{k=1}^d (I \otimes E_{ki})A(I \otimes E_{jk}) = B_{ij} \otimes I$, which implies that $B_{ij} \otimes I$ are in \mathcal{A} . So define \mathcal{B} as $\{B \in M(n, \mathbb{F}) : \text{such that } B \otimes I \in \mathcal{A}\}$, and we see that $\mathcal{A} = \mathcal{B}^{\{d\}}$. \square

Note that $(I \otimes M(d, \mathbb{F}))\mathcal{A}(I \otimes M(d, \mathbb{F})) = \mathcal{A}$ is equivalent to saying that \mathcal{A} is an $M(d, \mathbb{F})$ sub-bimodule of $M(n, \mathbb{F}) \otimes M(d, \mathbb{F})$, where we identify $M(d, \mathbb{F})$ with $I \otimes M(d, \mathbb{F})$. Similarly, for a subspace $W \leq \mathbb{F}^n \otimes \mathbb{F}^d$, one can see that W is of the form $W_0 \otimes \mathbb{F}^d$ if and only if $(I \otimes M(d, \mathbb{F}))W_0 = W_0$, that is, W_0 is an $M(d, \mathbb{F})$ -submodule of $\mathbb{F}^n \otimes \mathbb{F}^d$.

Shrunk subspaces in the blow-up situation.

Proposition 5.2. *If $\mathcal{A} = \mathcal{B}^{\{d\}}$ has an s -shrunk subspace, then \mathcal{A} has an s' -shrunk subspace where $s' \geq s$ such that d divides s' , and \mathcal{B} has an s'/d -shrunk subspace.*

Proof. As $\mathcal{A} = \mathcal{B}^{\{d\}}$, $(I \otimes M(d, \mathbb{F}))\mathcal{A}(I \otimes M(d, \mathbb{F})) = \mathcal{A}$. Assume that U is an s -shrunk subspace of \mathcal{A} : with $W = \mathcal{A}(U)$ we have $\dim_{\mathbb{F}} U - \dim_{\mathbb{F}} W = s$. Then $(I \otimes M(d, \mathbb{F}))W = (I \otimes M(d, \mathbb{F}))\mathcal{A}U = \mathcal{A}U = W$, thus $W = W_0 \otimes \mathbb{F}^d$ for some $W_0 \leq \mathbb{F}^n$. Similarly, as $\mathcal{A}(I \otimes M(d, \mathbb{F})) = \mathcal{A}$, we have $\mathcal{A}(I \otimes M(d, \mathbb{F}))U = W$, whence $U' = (I \otimes M(d, \mathbb{F}))U$ is an s' -shrunk subspaces with $s' \geq s$, and $U' = U_0 \otimes \mathbb{F}^d$ with some $U_0 \leq \mathbb{F}^n$. Note that $\dim(W)$, $\dim(U')$, and therefore s' , are all divisible by d . Noting $\mathcal{A} = \mathcal{B} \otimes M(d, \mathbb{F})$, we have $W_0 \leq \mathcal{B}(U_0)$ and so U_0 is an s'/d -shrunk subspace. \square

From the extension field to the original field. Assume that for some extension field \mathbb{K} of \mathbb{F} we are given a matrix $A' \in \mathcal{B} \otimes_{\mathbb{F}} \mathbb{K} \leq M(n, \mathbb{K})$ of rank r . Then, if $|\mathbb{F}| > r$, using the method of [18, Lemma 2.2], we can efficiently find a matrix $A \in \mathcal{B}$ of rank at least r . This procedure is also useful to keep sizes of the occurring field elements small. For completeness we include a brief description. Let $S \subseteq \mathbb{F}$ with $|S| = r + 1$ and let B_1, \dots, B_ℓ be an \mathbb{F} -basis for \mathcal{B} . Then $A' = a'_1 B_1 + \dots + a'_\ell B_\ell$, where $a'_i \in \mathbb{K}$. As A' is of rank r , there exists an $r \times r$ sub-matrix of A' with nonzero determinant. Assume that $a'_1 \notin S$. Then we consider the determinant of the corresponding sub-matrix of the polynomial matrix $x B_1 + a'_2 B_2 + \dots + a'_\ell B_\ell$. This determinant is a nonzero polynomial of degree at most r in x . Therefore there exists an element $a_1 \in S$ such that $a_1 B_1 + a'_2 B_2 + \dots + a'_\ell B_\ell$ has rank at least r . Continuing with a'_2, \dots, a'_ℓ , we can ensure that all the a_i 's are from S . Since the B_i 's span \mathcal{B} , the resulting matrix of rank at least r is in \mathcal{B} . We record this too as a fact.

Lemma 5.3 (Data reduction, [18, Lemma 2.2]). *Let $\mathcal{B} \leq M(k \times \ell, \mathbb{F})$ be given by a basis B_1, \dots, B_m , and let \mathbb{K} be an extension field of \mathbb{F} . Let S be a subset of \mathbb{F} of size at least $r + 1$. Suppose that we are given a matrix $A' = \sum a'_i B_i \in \mathcal{B} \otimes_{\mathbb{F}} \mathbb{K}$ of rank at least r . Then we can find $A = \sum a_i B_i \in \mathcal{B}$ of rank also at least r with $a_i \in S$. The algorithm uses $\text{poly}(k, \ell, r)$ rank computations for matrices of the form $\sum a''_i B_i$ where $a''_i \in \{a'_1, \dots, a'_m\} \cup S$.*

5.2 Regularity of blow-ups

Our goal in this subsection is to prove that when the field size is large enough, the maximum rank over $\mathcal{A} = \mathcal{B}^{\{d\}} \leq M(dn, \mathbb{F})$ is always divisible by d . The proof is constructive when $\text{char}(\mathbb{F}) = 0$, or when $\text{char}(\mathbb{F}) \nmid d$: if we get a matrix in \mathcal{A} of rank at least $rd + 1$, we will be able to construct a matrix in \mathcal{A} of rank at least $(r + 1)d$. This is the main technical tool to be used in the proof of Theorem 5.11.

We first present a version of the regularity lemma in which a matrix division algebra as in Lemma 4.5 is assumed to be part of the input.

Lemma 5.4 (Regularity of blow-ups, technical version). *Assume that we are given a matrix $A \in \mathcal{B}^{\{d\}} \leq M(dn, \mathbb{F})$ with $\text{rk}(A) = (r - 1)d + k$ for some $1 < k < d$. Let X and Y be formal variables and put $\mathbb{K} = \mathbb{F}'(X)$, where \mathbb{F}' is a finite extension of \mathbb{F} of degree at most d . Suppose further that $|\mathbb{F}| > (nd)^{\Omega(1)}$ and that we are also given a $\mathbb{K}(Y)$ -basis Γ of $M(d, \mathbb{K}(Y))$ such that the $\mathbb{K}(Y^d)$ -linear span of Γ is a central division algebra D' over $\mathbb{K}(Y^d)$. Let δ be the maximum of the degrees of the polynomials appearing as numerators or denominators of the entries of the matrices in Γ . Then, using $(nd + \delta)^{O(1)}$ arithmetic operations in \mathbb{F} , one can find a matrix $A'' \in \mathcal{B}^{\{d\}}$ with $\text{rk}(A'') \geq rd$. Furthermore, over \mathbb{Q} the bit complexity of the algorithm is polynomial in the size of the input data (that is, the total number of bits describing the entries of matrices and the coefficients of polynomials).*

Proof. Instead of the blow-up $\mathcal{B}^{\{d,d\}} = \mathcal{B} \otimes M(d, \mathbb{F})$ we start with the blow-up $\mathcal{B} \otimes M(d, \mathbb{K}(Y))$. Then both $\mathcal{B}^{\{d,d\}}$ and $\mathcal{B} \otimes_{\mathbb{F}} D'$ span, as a $\mathbb{K}(Y)$ -linear space, the blow-up $\mathcal{B} \otimes_{\mathbb{F}} M(d, \mathbb{K}(Y))$.

Claim 5.5. *Every matrix in $M(n, \mathbb{F}) \otimes D' \subset M(d, \mathbb{K}(Y))$ has rank (as a matrix over $\mathbb{K}(Y)$) divisible by d .*

Proof. Firstly note that $M(n, \mathbb{F}) \otimes_{\mathbb{F}} D' = M(n, \mathbb{K}(Z)) \otimes_{\mathbb{K}(Z)} D'$, since D' is a $\mathbb{K}(Z)$ -algebra. As $D' \subset M(d, \mathbb{K}(Y))$, $M(n, \mathbb{K}(Z)) \otimes_{\mathbb{K}(Z)} D'$ acts naturally on the $\mathbb{K}(Z)$ -space $\mathbb{K}(Z)^n \otimes \mathbb{K}(Y)^d \cong \mathbb{K}(Z)^n \otimes \mathbb{K}(Z)^{d^2} \cong \mathbb{K}(Z)^{nd^2}$. Since $D' \otimes_{\mathbb{K}(Z)} D'^{op} \cong M(d^2, \mathbb{K}(Z))$ [50, Corollary 15.5], it follows that the centralizer of this action is isomorphic to the opposite algebra D'^{op} . Therefore the image $A' \mathbb{K}(Y)^{dn}$ of any $A' \in M(n, \mathbb{F}) \otimes_{\mathbb{F}} D'$ is a D'^{op} -submodule, whence its dimension over $\mathbb{K}(Z)$ is divisible by d^2 . It follows that the dimension over $\mathbb{K}(Y)$ is divisible by d . \square

The claim enables us to “round up” the rank of A to the next multiple of d . Let B_1, \dots, B_d be an \mathbb{F} -basis of \mathcal{B} . Since $\text{rk}(A) > (r-1)d$ over \mathbb{F} , clearly $\text{rk}(A) > (r-1)d$ over $\mathbb{K}(Y)$ as well. Now Γ , the $\mathbb{K}(Z)$ -basis for D' is a $\mathbb{K}(Y)$ -basis for $M(d, \mathbb{K}(Y))$. Therefore A , as a matrix over $\mathbb{K}(Y)$, can be expressed as a linear combination (with coefficients over $\mathbb{K}(Y)$) of $\{B_i \otimes C : i \in [d], C \in \Gamma\}$. We use the method of Lemma 5.3 to find coefficients from $\mathbb{K}(Z)$ (or even from \mathbb{F}) such that the combination A' of the basis element for D' has rank also larger than $(r-1)d$. We have $A' \in \mathcal{B} \otimes D'$, whence by Claim 5.5, the rank of A' is at least rd . Then we express A' as a linear combination of elements – with coefficients from $\mathbb{K}(Y)$ – of an \mathbb{F} -basis of $\mathcal{B}^{\{d,d\}}$ which is also a $\mathbb{K}(Y)$ -basis for $\mathcal{B} \otimes M(d, \mathbb{K}(Y))$. Then we use again the algorithm of Lemma 5.3 to replace these coefficients to elements of \mathbb{F} to find a matrix $A'' \in \mathcal{B}$ of rank at least rd . \square

Using Remark 4.2 and Lemma 4.5 we immediately obtain the following results.

Lemma 5.6 (Regularity of blow-ups, non-constructive). *For $\mathcal{B} \leq M(n, \mathbb{F})$, assume that $|\mathbb{F}| > (nd)^{\Omega(1)}$. Then $\text{rk}(\mathcal{B}^{\{d\}})$ is divisible by d .*

Lemma 5.7 (Regularity of blow-ups, constructive). *For $\mathcal{B} \leq M(n, \mathbb{F})$ and $\mathcal{A} = \mathcal{B}^{\{d\}}$, assume that $\text{char}(\mathbb{F}) = 0$ or $\text{char}(\mathbb{F}) \nmid d$, and $|\mathbb{F}| > (nd)^{\Omega(1)}$. Then, given a matrix $A \in \mathcal{A}$ with $\text{rk}A > (r-1)d$, there exists a deterministic algorithm that returns $\tilde{A} \in \mathcal{A}$ of rank $\geq rd$. This algorithm uses $\text{poly}(nd)$ arithmetic operations and over \mathbb{Q} , all intermediate numbers have bit lengths polynomial in the input size.*

To see all the ingredients together, we give an expanded description in Algorithm 1.

Algorithm 1: Algorithm ROUNDUPRANK

Input: $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$. $S \subseteq \mathbb{F}$ with $|S| = (nd)^{\Omega(1)}$. $A \in \mathcal{B}^{\{d\}}$ of $\text{rk}(A) > (n-1)d$. If $\text{char}(\mathbb{F}) = p > 0$ then $p \nmid d$.

Output: $A'' \in \mathcal{B}^{\{d\}}$ with $\text{rk}(A'') = nd$.

- 1 $\mathbb{F}' \leftarrow \mathbb{F}[\zeta]$, where $\zeta \leftarrow$ a d th root of \mathbb{F} ;
- 2 $X_1, Y \leftarrow$ two independent formal variables;
- 3 $\Gamma = \{C_1, \dots, C_{d^2}\} \leftarrow$ a $\mathbb{F}'(X_1, Y)$ -basis of $M(d, \mathbb{F}'(X_1, Y))$ such that $\mathbb{F}'(X_1, Y^d)$ -linear span of Γ is a central division algebra over $\mathbb{F}'(X_1, Y^d)$;
- 4 Expand $A = \sum_{i \in [m], j \in [d^2]} \lambda_{i,j} B_i \otimes C_j$, $\lambda_{i,j} \in \mathbb{F}'(X_1, Y)$ as a matrix in $\mathcal{B} \otimes M(d, \mathbb{F}'(X_1, Y))$;
- 5 Compute $A' = \sum_{i \in [m], j \in [d^2]} \mu_{i,j} B_i \otimes C_j$, $\mu_{i,j} \in S$ such that $\text{rk}(A') = nd \geq \text{rk}(A)$;
- 6 $\Delta = \{E_{i,j} \mid i, j \in [d]\} \leftarrow$ the standard basis of $M(d, \mathbb{F}'(X_1, Y))$;
// That is $E_{i,j}(i, j) = 1$, and other entries are 0
- 7 Expand $A' = \sum_{i \in [m], j, k \in [d]} \nu_{i,j,k} B_i \otimes E_{j,k}$, $\nu_{i,j,k} \in \mathbb{F}'(X_1, Y)$;
- 8 Compute $A'' = \sum_{i \in [m], j, k \in [d]} \xi_{i,j,k} B_i \otimes E_{j,k}$, $\xi_{i,j,k} \in S$ such that $\text{rk}(A'') = \text{rk}(A')$;
- 9 **return** A'' .

We remind the reader how our previous development supports Algorithm 1. For Line 1, Section 4.4.1 describes how to work with $\mathbb{F}[\zeta]$. Line 3 calls the procedure in Lemma 4.5 which in turn utilizes Lemma 4.1. Line 4 and 7 are standard tasks in linear algebra. Line 5 and 8 rely on Lemma 5.3, which in turn calls Proposition 4.8. Lemma 4.6 ensures that the rank calculation is essentially correct.

We mention an immediate consequence.

Corollary 5.8. *Let $\mathcal{B} \leq M(n, \mathbb{F})$. Assume that the characteristic of \mathbb{F} is zero, $d \geq n$, and we are given a matrix $A \in \mathcal{B}^{\{d\}}$ with $\text{rk}A = rd$. Then, for any $d' > d$ we can efficiently find $\tilde{A} \in \mathcal{B}^{\{d'\}}$ with $\text{rk}\tilde{A} \geq rd'$.*

Proof. Induction on d' and d . Assume first that $d' = d + 1$. Then $\frac{d+1}{d} \leq \frac{n+1}{n} < \frac{r}{r-1}$, whence $(d+1)(r-1) < dr$. Therefore, if we embed $\mathcal{B}^{\{d\}}$ into $\mathcal{B}^{\{d+1\}}$ then A has rank $rd > (r-1)(d+1)$ and Lemma 5.7 applies. Proceed with $(d+2, d+1)$ in place of $(d+1, d)$. \square

We close this subsection with two open problems.

Remark 5.9.

1. It is an interesting problem to investigate whether (the non-constructive version of) the regularity lemma would hold over small fields. Giving a proof not relying on division algebras might be a progress in this direction.
2. Corollary 5.8 probably remains true for $d < n$. However, we do not see how to prove it.

5.3 Incrementing rank via blow-up

We have introduced in Section 2.2 a key technique here, namely the second Wong sequences. Recall that given $A \in \mathcal{B} \leq M(n, \mathbb{F})$, the second Wong sequence can be used to detect whether there exists a $\text{cork}(A)$ -shrunk subspace. If such a shrunk subspace exists, then $\text{ncrk}(\mathcal{B}) = \text{rk}(A)$. The difficulty is the case when such $\text{cork}(A)$ -shrunk subspace does not exist. One natural idea to proceed is to find $A' \in \mathcal{B}$ of rank $> \text{rk}(A)$, and use A' to test whether a $\text{cork}(A')$ -shrunk exists or not. If \mathcal{B} is rank-1 spanned, such an A' does exist, and another ingredient in [44] is an update procedure that finds A' of higher rank in this case. However, this in general is not possible, since $\text{rk}(\mathcal{B})$ and $\text{ncrk}(\mathcal{B})$ may differ. Fortunately, the concept of blow-ups helps: instead of looking for $A' \in \mathcal{B}$ of rank $> \text{rk}(A)$, we shall look for $A' \in \mathcal{B}^{\{d\}}$ of rank $> \text{rk}(A)d$ for some not too large $d \in \mathbb{N}$. It turns out that this is achievable, and an application of the regularity lemma even yields $A'' \in \mathcal{B}^{\{d\}}$ of rank $\geq (\text{rk}(A) + 1)d$.

Theorem 5.10. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ and $\mathcal{A} = \mathcal{B}^{\{d\}}$. Assume that we are given a matrix $A \in \mathcal{A}$ with $\text{rk}(A) = rd$. Let d' be an integer $> r$. Suppose that $|\mathbb{F}|$ is $(ndd')^{\Omega(1)}$, and if $\text{char}(\mathbb{F}) = p > 0$ then assume $p \nmid dd'$. There exists a deterministic algorithm that returns either an $(n-r)d$ -shrunk subspace for \mathcal{A} (equivalently, an $(n-r)$ -shrunk subspace for \mathcal{B}), or a matrix $A^* \in \mathcal{A} \otimes M(d', \mathbb{F})$ of rank at least $(r+1)dd'$. This algorithm uses $\text{poly}(ndd')$ arithmetic operations and, over \mathbb{Q} , all intermediate numbers have bit lengths polynomial in the input size.*

Proof. We present the algorithm formally as Algorithm 2.

Let us first outline what the algorithm does. From Line 1 to 6 it computes the second Wong sequence with respect to (A, \mathcal{A}) . Line 7 and 8 deal with the case when the sequence provides a $\text{cork}(A)$ -shrunk subspace. In the other case, we first utilize the sequence to get a matrix A'' of rank $> rdd'$ (Line 9 to 25). Then we obtain the desired A^* , by applying the regularity lemma (Lemma 5.7) to A'' .

We now explain some implementation details of the algorithm.

Line 3, $\ell \leq r+1$ The second Wong sequence, when applied to matrix spaces of the form $\mathcal{A} = \mathcal{B}^{\{d\}}$, stabilizes faster because of the following. Since $(I \otimes M(d, \mathbb{F}))\mathcal{A} = \mathcal{A}$, at stage j we have $(I \otimes M(d, \mathbb{F}))\mathcal{A}W_j = \mathcal{A}W_j$, whence the dimension of $\mathcal{A}W_j$ is divisible by d for every j . It follows that, until stabilization, the dimension of $\mathcal{A}W_j$ increases by at least d and so the sequence stabilizes to its limit in at most $r+1$ steps when applied to $A \in \mathcal{A}$ of rank rd .

Line 9 To compute C_i 's, one perform the following. Take a basis of \mathcal{A} . Search for a basis element Y such that $YA^{-1}(\mathcal{A}A^{-1})^{\ell-1}(0) \not\subseteq \text{im}(A)$. Put $C_\ell = Y$ and search for a basis element Y such that $C_\ell A^{-1}YA^{-1}(\mathcal{A}A^{-1})^{\ell-2}(0) \not\subseteq \text{im}(A)$. Continue this iteration and the desired C_i 's can be computed.

Algorithm 2: Algorithm INCREMENTRANK

Input: $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$. $A \in \mathcal{A} = \mathcal{B}^{\{d\}}$ of $\text{rk}(A) = rd$. An integer $d' > r$ such that if $\text{char}(\mathbb{F}) = p > 0$ then $p \nmid d'$. $S \subseteq \mathbb{F}$ with $|S| = (ndd')^{\Omega(1)}$.

Output: One of the following: (1) An $(n-r)d$ -shrunk subspace of \mathcal{A} ; (2) $A^* \in \mathcal{A}^{\{d'\}} = \mathcal{B}^{\{dd'\}}$ with $\text{rk}(A^*) \geq (r+1)dd'$.

```

1  $W_0, W^* \leftarrow \mathbf{0}$ ;
2  $\ell = 0$ ;
3 while  $W^* \subseteq \text{im}(A)$  and  $\ell \leq r+1$  do
4    $\ell \leftarrow \ell + 1$ ;
5    $W_\ell \leftarrow \mathcal{A}A^{-1}(W_{\ell-1})$ ;
6    $W^* = W_\ell$ .
7 if  $W^* \subseteq \text{im}(A)$  then
8   return  $A^{-1}(W^*)$ .
// In the following  $\ell$  is the smallest integer  $i$  such that  $W_i \not\subseteq \text{im}(A)$ .
9 Compute  $C_1, \dots, C_\ell \in \mathcal{A}$  such that  $C_\ell A^{-1} C_{\ell-1} A^{-1} \dots C_1 A^{-1}(0) \not\subseteq \text{im}(A)$ ;
10 Take  $v_1 \in A^{-1}(0) = \ker(A)$  such that  $C_\ell A^{-1} C_{\ell-1} A^{-1} \dots C_1(v_1) \notin \text{im}(A)$ ;
11 for  $i = 2$  to  $\ell$  do
12    $v_i \leftarrow A^{-1} C_{i-1}(v_{i-1})$ .
// Note that  $C_\ell v_\ell \notin \text{im}(A)$ .
13  $\{u_1, \dots, u_{d'}\} \leftarrow$  a basis of  $\mathbb{F}^{d'}$ ;
14 for  $i = 1$  to  $\ell$  do
15   if  $i = r+1 = d'$  then
16      $Z_i \leftarrow$  the matrix in  $M(d', \mathbb{F})$  such that  $Z_{r+1}(u_{r+1}) = u_1$  and  $Z_{r+1}(u_j) = 0$  for  $j \neq r+1$ .
17   else
18      $Z_i \leftarrow$  the matrix in  $M(d', \mathbb{F})$  such that  $Z_i(u_i) = u_{i+1}$  and  $Z_i(u_j) = 0$  for  $j \neq i$ .
19  $A' \leftarrow A \otimes I$ ;
20  $C' \leftarrow C_1 \otimes Z_1 + \dots + C_\ell \otimes Z_\ell$ ;
21 if  $\text{rk}(C') > rdd'$  then
22    $A'' \leftarrow C'$ .
23 else
24   Compute  $\lambda \in S$  such that  $A' + \lambda C'$  is of  $> rdd'$ ;
25    $A'' \leftarrow A' + \lambda C'$ .
26 Compute  $A^*$  of rank  $\geq (r+1)dd'$  using  $A''$  and Lemma 5.7;
27 return  $A^*$ .

```

That Algorithm 2 runs in the stated time bound follows easily from Fact 2.1 and Lemma 5.7. To see the correctness of the algorithm, it remains to prove that in Line 21 to 25 we do obtain A'' of rank $> rdd'$. Consider the vectors $w_1 = v_1 \otimes u_1, \dots, w_t = v_t \otimes u_t$. We now observe that: (1) $w_1 \in \ker A'$; (2) $A'w_j = C'w_{j-1}$ for $j = 2, \dots, t$; (3) $C'w_t = C_t v_t \otimes u_{t+1} \notin (\mathcal{AF}^{nd}) \otimes \mathbb{F}^{d'}$; as $\mathcal{AF}^{nd} \otimes \mathbb{F}^{d'} \supseteq A' \mathbb{F}^{ndd'}$, we have $C'w_t \notin A' \mathbb{F}^{ndd'}$. This means that the limit of the second Wong sequence for the pair $(A', \langle A', C' \rangle)$ runs out of the image of A' . By Fact 2.2, A' is not of maximum rank in $\langle A', C' \rangle$, and Line 21 to 25 just describe a straightforward method to obtain a matrix of highest rank in a 2-dimensional matrix space. \square

An iteration based on Theorem 5.10 proves the following Theorem 5.11. Note that in Theorem 5.10, d' can be chosen as either $r + 1$ or $r + 2$, depending on which is not divisible by $\text{char}(\mathbb{F})$.

Theorem 5.11. *Suppose we are given $\mathcal{B} := \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, and $A \in \mathcal{B}$ with $\text{rk}(A) = s < n$. Let $d = (n+1)!/(s+1)!$, and assume that $|\mathbb{F}| = \Omega(nd)$. Then there exists a deterministic algorithm, that computes a matrix $B \in \mathcal{B} \otimes M(d', \mathbb{F})$ of rank rd' for some $d' \leq d$ and, if $r < n$, an $(n-r)$ -shrunk subspace for \mathcal{B} . The algorithm uses $\text{poly}(n, d)$ arithmetic operations, and when working over \mathbb{Q} , has bit complexity polynomial in n, d and the input size.*

Now Corollary 1.5 and 1.6 follow easily. To see Corollary 1.5, note that if we choose a matrix in \mathcal{B} randomly, it will be of maximum rank. Using that matrix as A in Theorem 5.11, Corollary 1.5 is proved. For Corollary 1.6, if \mathcal{B} has no shrunk subspace, a full-rank matrix will be certainly present in $M(d', \mathbb{F}) \otimes \mathcal{B}$ for some $d' \leq (n+1)!$, giving us the upper bound on $\sigma(R(n, m))$.

Acknowledgements. We would like to thank Mátyás Domokos, Bharat Adsul, Ketan Mulmuley, Partha Mukhopadhyay and K. N. Raghavan for discussions related to this work. Part of the work was done when Youming was visiting the Simons Institute for the program Algorithms and Complexity in Algebraic Geometry and when Gábor was visiting the Centre for Quantum Technologies, National University of Singapore. Research of the first author was also supported in part by the Hungarian National Research, Development and Innovation Office – NKFIH, Grants NK105645 and 115288. Youming’s research was supported by the Australian Research Council DECRA DE150100720.

References

- [1] B. Adsul, S. Nayak, and K. V. Subrahmanyam, *A geometric approach to the Kronecker problem II: rectangular shapes, invariants of matrices and the Artin–Procesi theorem*, preprint, 2007.
- [2] Bharat Adsul and K. V. Subrahmanyam, *A geometric approach to the Kronecker problem I: the two row case*, Proceedings Mathematical Sciences **118** (2008), no. 2, 213–226.
- [3] S.A Amitsur, *Rational identities and applications to algebra and geometry*, Journal of Algebra **3** (1966), no. 3, 304 – 359.
- [4] MD Atkinson and S Lloyd, *Primitive spaces of matrices of bounded rank*, Journal of the Australian Mathematical Society (Series A) **30** (1981), no. 04, 473–482.
- [5] George W. Bergman, *Skew fields of noncommutative rational functions (preliminary version)*, Sminaire Schtzenberger **1** (1969-1970), 1–18 (eng).
- [6] Peter A. Brooksbank and Eugene M. Luks, *Testing isomorphism of modules*, Journal of Algebra **320** (2008), no. 11, 4020–4029.
- [7] M. Bürgin and J. Draisma, *The Hilbert null-cone on tuples of matrices and bilinear forms*, Mathematische Zeitschrift **254** (2006), no. 4, 785–809.
- [8] Peter Bürgisser, J. M. Landsberg, Laurent Manivel, and Jerzy Weyman, *An overview of mathematical issues arising in the geometric complexity theory approach to $VP \neq VNP$* , SIAM J. Comput. **40** (2011), no. 4, 1179–1209.

- [9] Jonathan F. Buss, Gudmund S. Frandsen, and Jeffrey O. Shallit, *The computational complexity of some problems of linear algebra*, J. Comput. Syst. Sci. **58** (1999), no. 3, 572–596.
- [10] Marco Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova, *Tighter connections between derandomization and circuit lower bounds*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA, 2015, pp. 645–658.
- [11] Alexander L. Chistov, Gábor Ivanyos, and Marek Karpinski, *Polynomial time algorithms for modules over finite dimensional algebras*, ISSAC, 1997, pp. 68–74.
- [12] Ajeh M Cohen, Gábor Ivanyos, and David B Wales, *Finding the radical of an algebra of linear transformations*, Journal of Pure and Applied Algebra **117** (1997), 177–193.
- [13] P. M. Cohn, *The word problem for free fields*, J. Symbolic Logic **38** (1973), no. 2, 309–314.
- [14] ———, *The word problem for free fields: A correction and an addendum*, J. Symbolic Logic **40** (1975), no. 1, 69–74.
- [15] ———, *Free rings and their relations*, L.M.S. Monographs, Acad. Press, 1985, First edition 1971.
- [16] ———, *Skew fields: Theory of general division rings*, Encyclopedia of Mathematics and its Applications, Cambridge University Press, 1995.
- [17] P. M. Cohn and C. Reutenauer, *On the construction of the free field*, International Journal of Algebra and Computation **9** (1999), no. 3-4, 307–323.
- [18] Willem A. de Graaf, Gábor Ivanyos, and Lajos Rónyai, *Computing Cartan subalgebras of Lie algebras*, Applicable Algebra in Engineering, Communication and Computing **7** (1996), no. 5, 339–349.
- [19] Harm Derksen, *Polynomial bounds for rings of invariants*, Proceedings of the American Mathematical Society **129** (2001), no. 4, 955–964.
- [20] Harm Derksen and Visu Makam, *Polynomial degree bounds for matrix semi-invariants*, preprint ArXiv:1512.03393, 2015.
- [21] Harm Derksen and Jerzy Weyman, *Semi-invariants of quivers and saturation for littlewood-richardson coefficients*, Journal of the American Mathematical Society **13** (2000), no. 3, 467–479.
- [22] M. Domokos, *Poincaré series of semi-invariants of 2×2 matrices*, Linear Algebra and its Applications **310** (2000), no. 1, 183–194.
- [23] ———, *Relative invariants of 3×3 matrix triples*, Linear and Multilinear Algebra **47** (2000), no. 2, 175–190.
- [24] ———, *Finite generating system of matrix invariants*, Math. Pannon **13** (2002), no. 2, 175–181.
- [25] M. Domokos and V. Drensky, *Defining relation for semi-invariants of three by three matrix triples*, Journal of Pure and Applied Algebra **216** (2012), no. 10, 2098–2105.
- [26] M. Domokos, S. G. Kuzmin, and A. N. Zubkov, *Rings of matrix invariants in positive characteristic*, Journal of Pure and Applied Algebra **176** (2002), no. 1, 61–80.
- [27] M. Domokos and A. N. Zubkov, *Semi-invariants of quivers as determinants*, Transformation groups **6** (2001), no. 1, 9–24.
- [28] Stephen Donkin, *Invariants of several matrices*, Inventiones mathematicae **110** (1992), no. 1, 389–401.

[29] Stephen Donkin, *Invariant functions on matrices*, Mathematical Proceedings of the Cambridge Philosophical Society **113** (1993), 23–43.

[30] Jack Edmonds, *Systems of distinct representatives and linear algebra*, J. Res. Nat. Bur. Standards Sect. B **71** (1967), 241–245.

[31] David Eisenbud and Joe Harris, *Vector spaces of matrices of low rank*, Advances in Mathematics **70** (1988), no. 2, 135 – 155.

[32] Edward Formanek, *Generating the ring of matrix invariants*, Ring Theory (Freddy M. J. van Oystaeyen, ed.), Lecture Notes in Mathematics, vol. 1197, Springer Berlin Heidelberg, 1986, pp. 73–82 (English).

[33] M. Fortin and C. Reutenauer, *Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank*, Séminaire Lotharingien de Combinatoire **52** (2004), B52f.

[34] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson, *A deterministic polynomial time algorithm for non-commutative rational identity testing*, preprint ArXiv:1511.03730, 2015.

[35] James Geelen and Satoru Iwata, *Matroid matching via mixed skew-symmetric matrices*, Combinatorica **25** (2005), no. 2, 187–215.

[36] James F. Geelen, *Maximum rank matrix completion*, Linear Algebra and its Applications **288** (1999), 211–217.

[37] James F Geelen, *An algebraic matching algorithm*, Combinatorica **20** (2000), no. 1, 61–70.

[38] James F. Geelen, Satoru Iwata, and Kazuo Murota, *The linear delta-matroid parity problem*, Journal of Combinatorial Theory, Series B **88** (2003), no. 2, 377–398.

[39] Leonid Gurvits, *Classical complexity and quantum entanglement*, J. Comput. Syst. Sci. **69** (2004), no. 3, 448–484.

[40] Leonid Gurvits and Peter N. Yianilos, *The deflation-inflation method for certain semidefinite programming and maximum determinant completion problems (extended abstract)*, Tech. report, NECI, 1998.

[41] Nicholas J. A. Harvey, David R. Karger, and Kazuo Murota, *Deterministic network coding by matrix completion*, Proceedings of SODA, ACM-SIAM, 2005, pp. 489–498.

[42] D. Hilbert, *Über die vollen invariantensysteme*, Math. Ann. (1893), no. 42, 313–370.

[43] Pavel Hrubeš and Avi Wigderson, *Non-commutative arithmetic circuits with division*, Theory of Computing **11** (2015), 357–393.

[44] Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha, *Generalized wong sequences and their applications to edmonds' problems*, J. Comput. Syst. Sci. **81** (2015), no. 7, 1373–1386.

[45] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena, *Deterministic polynomial time algorithms for matrix completion problems*, SIAM J. Comput. **39** (2010), no. 8, 3736–3751.

[46] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam, *Constructive non commutative rank computation in deterministic polynomial time over fields of arbitrary characteristics*, preprint arXiv:1512.03531, 2015.

[47] ———, *On generating the ring of matrix semi-invariants*, preprint, 2015.

[48] Valentine Kabanets and Russell Impagliazzo, *Derandomizing polynomial identity tests means proving circuit lower bounds*, Computational Complexity **13** (2004), no. 1-2, 1–46.

[49] Erich Kaltofen, *On computing determinants of matrices without divisions*, Proceedings of the 1992 International Symposium on Symbolic and Algebraic Computation, ISSAC '92, Berkeley, CA, USA, July 27-29, 1992, 1992, pp. 342–349.

[50] T.Y. Lam, *A first course in noncommutative rings*, Graduate Texts in Mathematics, Springer, 1991.

[51] Nathan Linial, Alex Samorodnitsky, and Avi Wigderson, *A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents*, Combinatorica **20** (2000), no. 4, 545–568.

[52] László Lovász, *Singular spaces of matrices and their application in combinatorics*, Boletim da Sociedade Brasileira de Matemática-Bulletin/Brazilian Mathematical Society **20** (1989), no. 1, 87–99.

[53] Meena Mahajan and V. Vinay, *Determinant: Combinatorics, algorithms, and complexity*, Chicago Journal of Theoretical Computer Science **1997** (1997), no. 5.

[54] Peter Malcolmson, *A prime matrix ideal yields a skew field*, Journal of the London Mathematical Society **s2-18** (1978), no. 2, 221–233.

[55] Laurent Manivel, *A note on certain Kronecker coefficients*, Proceedings of the American Mathematical Society **138** (2010), no. 1, 1–7.

[56] Ketan Mulmuley, *A fast parallel algorithm to compute the rank of a matrix over an arbitrary field*, Combinatorica **7** (1987), no. 1, 101–104.

[57] ———, *On P vs. NP and geometric complexity theory: Dedicated to Sri Ramakrishna*, J. ACM **58** (2011), no. 2, 5.

[58] ———, *Geometric complexity theory V: equivalence between blackbox derandomization of polynomial identity testing and derandomization of noether's normalization lemma*, 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, 2012, pp. 629–638.

[59] Ketan Mulmuley and Milind A. Sohoni, *Geometric complexity theory I: an approach to the P vs. NP and related problems*, SIAM J. Comput. **31** (2001), no. 2, 496–526.

[60] ———, *Geometric complexity theory II: towards explicit obstructions for embeddings among class varieties*, SIAM J. Comput. **38** (2008), no. 3, 1175–1206.

[61] Kazuo Murota, *Matrices and matroids for systems analysis*, Springer, 2000.

[62] Vladimir L Popov, *The constructive theory of invariants*, Izvestiya: Mathematics **19** (1982), no. 2, 359–376.

[63] C. Procesi, *The invariant theory of $n \times n$ matrices*, Advances in Mathematics **19** (1976), no. 3, 306–381.

[64] Ju. P. Razmyslov, *Trace identities of full matrix algebras over a field of characteristic zero*, Mathematics of the USSR-Izvestiya **8** (1974), no. 4, 727, English translation available at <http://iopscience.iop.org/0025-5726/8/4/A01>.

[65] T. G. Room, *The geometry of determinantal loci*, The Cambridge University Press, 1938.

[66] Aidan Schofield and Michel Van den Bergh, *Semi-invariants of quivers for arbitrary dimension vectors*, Indagationes Mathematicae **12** (2001), no. 1, 125–138.

[67] W. T. Tutte, *The factorization of linear graphs*, Journal of the London Mathematical Society **1** (1947), no. 2, 107–111.

- [68] E. Witt, *Zyklische Körper und Algebren der Charakteristik p vom Grad pn . Struktur diskret bewerteter perfekter Körper mit vollkommenem Restklassenkörper der Charakteristik p* , J. Reine Angew. Math, **176** (1937), no. 01, 126–140.
- [69] Kai-Tak Wong, *The eigenvalue problem $\lambda Tx + Sx$* , Journal of Differential Equations **16** (1974), no. 2, 270 – 280.

A Derksen's bound applied to $R(n, m)$

Proof. We just need to indicate certain parameters for the matrix semi-invariants that are used in Derksen's bound.

Suppose a group G acts on a vector space V rationally, and let R be the resulting invariant ring. Theorem 1.1 in [19] shows that the degree bound is upper bounded by $\max(2, 3/8 \cdot s \cdot \sigma(R)^2)$, where σ is the degree bound for defining the nullcone, and s is the dimension of R .

s is upper bounded by the number of variables. Therefore for $R(m, n)$, $s \leq mn^2 \leq n^4$.

To bound σ , we use Proposition 1.2 in [19]. Recall that G as an algebraic group, is defined by a system of polynomial equations in z_1, \dots, z_t . For example, $\mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ is defined by $\det(X) = 1$ and $\det(Y) = 1$, where X and Y are $n \times n$ variable matrices. The action of G is rational, so it can be recorded as $\rho: G \rightarrow \mathrm{GL}(V)$ by $g \rightarrow (a_{i,j}(g))_{i,j \in \mathrm{dim}(V)}$, where each $a_{i,j}$ is a polynomial in z_1, \dots, z_t .

$\sigma(R)$ is then upper bounded by $H^{t-m} A^m$, where t is the number of variables used to define G as above, $m = \mathrm{dim}(G)$, H is the maximum degree over polynomials defining G , and A is the maximum degree over polynomials defining the action. So for $R(n, m)$, $t = 2n^2$, $m = 2n^2 - 2$, $H = n$, and $A = 2$. It follows that $\sigma(R(n, m)) \leq n^2 \cdot 2^{2n^2-2}$.

Therefore $R(n, m)$ is generated by elements of degree $\leq 3/128 \cdot n^8 \cdot 16^{n^2}$. □